



The Palgrave Handbook of Technological Finance

Edited by
Raghavendra Rau
Robert Wardrop
Luigi Zingales

palgrave
macmillan

The Palgrave Handbook of Technological Finance

Raghavendra Rau · Robert Wardrop ·
Luigi Zingales
Editors

The Palgrave
Handbook
of Technological
Finance

palgrave
macmillan

Editors

Raghavendra Rau
University of Cambridge
Cambridge, UK

Robert Wardrop
University of Cambridge
Cambridge, UK

Luigi Zingales
University of Chicago Booth School of
Business
University of Chicago
Chicago, IL, USA

ISBN 978-3-030-65116-9 ISBN 978-3-030-65117-6 (eBook)
<https://doi.org/10.1007/978-3-030-65117-6>

© The Editor(s) (if applicable) and The Author(s) 2021

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Cover credit: MarsYu

This Palgrave Macmillan imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Finance has always taken advantage of technological innovation. No sooner had the telegraph been invented than traders were using it to arbitrage the price of stocks across cities. With the telephone, came the first forms of home banking. The introduction of computers transformed stock exchanges from overcrowded and loud pits to silent rooms filled with microprocessors. Yet, for most of the last two millennia, technological innovation has only allowed finance to do roughly what had been done before, though it has allowed us to do the same things better, faster, and more cheaply. The loans granted in the 1990s were not so different from the loans described in the Hammurabi code (1754 BCE) nor was it different the way those loans were granted almost four millennia later.

To be sure, technological innovation in finance has allowed financial transactions to be carried out at scale. At the end of the last decade of the twentieth century, for example, enormous amounts of money were wired daily around the world, credit and charge cards became the standard form of payment around most of the developed world, and credit started to be granted with credit score models. Still, international wire transfers were just a faster version of the letters of credit used during the Crusades. Credit card payments were a more sophisticated version of store credit, present since stores were invented. Even the most advanced credit score models resemble very closely the way yesteryear bankers were granting credit: they use the same qualitative and quantitative variables. In other words, technology reduced the cost of collecting and processing information, but it did not substantially change the business

of making payments and loans nor did it eliminate the essential frictions these involved, from adverse selection to moral hazard. A loan at the time of Hammurabi faced roughly the same moral hazard and adverse selection risk of a loan at the time of George Bush.

Not only did technological progress in the past not disrupt the way finance was done, it did not disrupt the existing financial institutions, which were the only institutions with the scale and resources to apply these innovations first. Thanks to the telegraph, stock exchanges became more efficient. Thanks to automated tellers, incumbent banks were able to reach every corner of the earth. Thanks to computers, large trading houses could reduce the cost of trading to a fraction of what it used to be.

At the beginning of the third millennium, however, the confluence of three major innovations has started to revolutionize the way finance is conducted. The first is the widespread introduction of sensor technology beginning with the first smartphone, the iPhone, in 2007. Thanks to smartphones, today's average individual has at her fingertips thousand times more computing power than was necessary to send man to the moon, more information than the best library used to contain, and more communicating power than any propaganda machine ever dreamed of possessing. More importantly, smartphones convinced most human beings to wear tracking devices, once reserved only for convicted felons on parole. Not only is it now possible to know whom we have talked to, but also where we have been, near whom, and for how long. Smartphones can track what searches we carried out, what books we bought, what vacations we shopped for, what dreams we have, and which diseases we fear to have. The latest iterations of smart devices add biometric information to the mix. These allow phones to capture additional information that we may not know ourselves. The smartphone knows when you are stressed, when you are low on sugar, and when you like a person of the same or opposite sex. In so doing, smartphones have enabled a degree of constant surveillance, a panopticon, that even Bentham or Orwell would have struggled to conceive. While this surveillance poses very serious political problems, it does create enormous opportunities to eliminate the frictions of financing.

These frictions are largely related to the asymmetry of information. Adverse selection and moral hazard problems are intrinsically linked to the inability to observe some individual characteristics or some actions, respectively. In fact, adverse selection used to be called 'hidden information' and moral hazard 'hidden actions'. If neither information nor actions can be hidden, the financing frictions related to the asymmetry of information, which have dominated the finance literature for the better part of the second half of the twentieth century, are eliminated.

In addition, smartphones have also reduced, if not eliminated the cost of coordination. All human exchange requires some degree of coordination. A buyer needs to find a willing seller and vice versa. From the Roman Forum to the Middle Age fair, this coordination has been resolved through the creation of physical places where people could meet. Geographical distance, however, constituted a major impediment to the liquidity of those markets, as it did the cost of spending the better part of a working day waiting for the counterparty to show up.

To facilitate the matching of buyers and sellers and create liquidity, financial institutions also created standardized contracts. Commodity futures contracts, for example, carefully specify what has to be delivered, when it should be delivered, and where it should be delivered. If the contracts on offer do not match the buyer's specific needs, the buyer has to incur basic risk.

Smartphones have solved all these problems. Buyers and sellers do not need to meet in the same physical location, since they can easily find each other in cyberspace, as eBay or Taobao has shown. Buyers and sellers do not need to spend time waiting for each other, because thanks to the internet, they can arrange to meet virtually at pre-determined time in a virtual space. Facilitating liquidity, these innovations have reduced the need for standardization. If a buyer can find a seller matched along all the dimensions she values, she does not need to converge to the standardized liquid market, but she can trade in a bespoke one.

On the one hand, this flexibility has enormously expanded the options for individuals. On the other hand, it has led to privatization of many of these markets. When a physical location was involved, the market itself was either publicly owned or jointly owned through a cooperative. In contrast, virtual markets are generally owned by a private entity, which maximizes its profits, not necessarily the quantity traded.

Finally, smartphones have radically changed the customer experience. Historically, merchants have chosen what to show to a customer and how to show it. The only limitations were imposed by regulation and/or competition. But competition only worked inasmuch as customers interacted repeatedly with the same firms. In most consumer financial transactions, however, customers participate in the market relatively infrequently (for example, when buying a house) and sometimes learn about their mistakes only many decades later (for example, when they buy insurance, or when they invest in pension funds). The computing power of the smartphone, however, can dramatically change this situation. Information can easily be collected and rearranged in a customer-friendly way. Even legal clauses can easily be scanned by an

algorithm and be flagged out as potentially dangerous for the uninformed customer.

Today, smartphones have been augmented by the Internet of things which comes with near-ubiquitous coverage by a huge number of sensors scattered in our houses or in our neighbourhoods that gather an even larger amount of information on us, in many cases, without us even being aware of the extent of the information being gathered.

While the introduction of widespread sensor technology is by far the most important innovation, its applications to financial services would be relatively limited if it were not for parallel developments in artificial intelligence (AI) and cryptography. Most human beings do not have either the time or the ability to interpret the massive amount of data produced by smartphones and sensors. Artificial intelligence has been defined as the 'ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation' (Kaplan and Haenlein (2019)). Thus, the application of AI to the massive amount of data produced by sensors changes the way financial services are offered and the way they are purchased. The use of AI need not be benign. It is possible to imagine your smartphone monitoring your sugar levels and sending you an advertisement for an impulse buy of clothes when your sugar levels are low and you have little ability to resist.

Transacting at a distance, without trusted intermediaries, would be too costly without the modern developments in cryptography. Currently, securely transferring payment and other confidential information requires the use of a trusted intermediary, usually a bank, who extracts large rents for this certification role. Commodity trade finance, for example, is a complicated process that involves exchanging signed original paper letters of credit that require frequent amendments and involve participants from around the world who do not know each other. By putting the verification process onto an immutable encrypted blockchain allows buyers and sellers to transfer information with no possibility of tampering, and most importantly, no central rent-extracting intermediary.

In sum, the combination of sensors, AI, and cryptography is revolutionizing the way finance is provided. As with all revolutions, the fintech one is likely to bring about a lot of improvements but will also leave a lot of blood in the streets. In this handbook, we try to provide a comprehensive analysis of the way in which the provision of financial services is changing, through a series of contributions by leading experts. We also try to provide a reasoned guide of how this evolution will change the financial landscape. By its very nature, this is a speculative exercise, since revolutions often take unexpected

directions. Nevertheless, we think it is a necessary step towards an understanding of the changes that are taking place under our eyes. We are finishing this handbook as governments around the world respond to the COVID-19 virus with policies that have dramatically slowed economic activity. While no one can predict with certainty what the post-COVID financial services landscape will look like, evidence suggests that the rate of adoption of digital financial services has been accelerating in many countries as a result of the crisis. As a result, regulators and policy-makers may be confronted with many of the issues raised in this handbook much sooner than we would have anticipated even a year ago.

In analysing these changes, we will distinguish to what extent they are opportunistic or really value-enhancing. In some instances, the opportunistic nature of the changes is clear, as when technology is applied to arbitrage regulatory requirements. In others, it is ambiguous. If applied by all lenders, a better way to screen borrowers helps allocate credit where it is most needed. Yet, when one lender introduces a better credit-scoring model, she also benefits by being able to dump hidden bad credit risks to her competitors. For example, a US Midwest insurance firm, Root insurance, uses detailed smart-phone location and acceleration data to target only careful drivers. Other insurance companies in the area not using this technology are stuck with poor drivers. Similarly, using AI to make credit allocation decisions, might lead to discrimination on racial, gender, or geographic grounds.

Technology is not intrinsically good or bad. The same computational power that enables banks to maximize the number of overdraft fees charged to their poor customers also empowers new companies to offer a management system to their customers for free to avoid the very same charges. What makes technology good or bad is the use we make of it. Competition and regulation determine the way technology is used and, thus, whether it is used to the benefit of humankind or its detriment.

As economists, we should know that competition plays a crucial role in ensuring that technology is used for the benefit of everyone, yet all too often we forget this truism. With all its benefits, technological innovation carries also a major risk: a tendency towards monopolies. The increasing return to scale provided by data, the close to zero marginal costs of expanding a digital business around the world, and the network externalities present in many of the finance functions, like payment, create the conditions for a winner-take-all market. In the United States, the entrenched position of existing players makes this tendency less visible, but in other countries, where the incumbents are weak or non-existent, it is blatant. Kenya provides a striking example. While M-Pesa created an efficient way to transact at a distance, providing

access to financial services to a mostly unbanked population, it did so by creating a monopoly in the payment market and conferring enormous market power to one mobile communication player. As a result, Kenyans have to pay hefty margins for both services. In China, where financial institutions were not very developed, Ant Group is creating a conglomerate with a dominant position in all services, from payments to credit, from wealth management to insurance. Its power is not limited to China. Ant Group has major stakes in mobile payment systems the world over: from PayTM in India to PayGo in Indonesia, from Kakao Bank in Korea to MercadoLibre in Argentina.

Network externalities and increasing return to scale, however, are not just immutable technological features, they are also the product of regulatory choices. If phone companies do not enjoy network externalities, it is because regulation forces all phone companies to accept calls coming from the clients of other companies. In the same way, the reason why we can easily switch from a mobile provider to another, carrying our phone number with us, is because in most developed countries, regulation has forced companies to accept number portability.

The same is true in most finance markets. Mandatory interoperability could resolve the monopoly problem in the Kenyan payment system. The problem is that interoperability needs to be not only mandated, but also enforced, which is not always easy when a lot of technical steps are involved. Interoperability is better achieved by creating a common standard through a public–private partnership, like the Unified Payment Interface developed in India by the National Payments Corporation, a not for profit enterprise owned by the Reserve Bank of India and a consortium of major banks.

As network externalities can be eliminated, so can the barrier to entry created by data ownership. If the ultimate ownership of the data is allocated to customers (as done in Europe), customers can easily transfer their data across banks to create more competition for their business. This is the purpose of the European Open Banking directive, which mandates an open Application Process Interface (API) to enable clients to transfer data across banks. While at the time of writing, the practical success of this initiative is still in question, it clearly goes in the right direction.

In the world of finance, especially household finance, competition—while necessary—is not sufficient to ensure that technology will bring benefits to everyone. This is where regulation plays a crucial role: regulators are entrusted with the very difficult role of separating the social value of innovation from its market value (i.e., its ability to generate profits). Faced with this challenge, it is easy for regulators to fall into one of the two extremes. One extreme is the ‘just say no’ approach. After all, no regulator loses her job for failing

to license a new startup. At the other extreme, there is the techno-optimist regulator who, inspired by the potential of the technology or by a lucrative future career in the industry, considers all profitable innovations as welfare enhancing. For this reason, it is especially important that regulators understand not only how technology is changing finance, but also what the welfare effects of this innovation are.

Last but not least, financial regulators will share some of the responsibility in promoting competition. Historically, financial regulators were more concerned with financial stability than with competition, often happy to sacrifice the latter on the altar of the former. Today, if they want technological improvements to benefit consumers, financial regulators have to ensure a lively competition. This goal cannot be delegated uniquely to the antitrust authorities. When it comes to creating common standards, financial regulators have to play a major role as played by the Reserve Bank of India for the Indian payment system. Performing this role, however, requires regulators to have the technical skills to understand how competition can be promoted in a rapidly evolving field. We hope this handbook provides them with a useful guide in this sense.

Each of the six sections in the handbook in the handbook address an aspect of technological finance of interest to regulators in the context of regulatory objectives and impact on society. The thirty-one chapters have been authored by a group of academics and leading expert practitioners from around the world with a deep understanding of developments in technological finance and the implication for regulation and policy-making. The geographic diversity of the authors reflects the transnational nature of technological finance and the implication for the global regulatory and policy-making community. In short, doing nothing and not responding to these developments is not a viable option for regulators because whatever innovation emerges in one regulatory jurisdiction is likely to propagate and appear in many others.

The first section of this handbook includes two chapters providing an introduction to technological finance by examining its impact on coordination and personalization in both present and historical terms. The chapter from Lin and Brown focus on the present by providing an overview of the benefits and risks to consumers posed by new channels and instruments of finance enabled by technology-enabled innovation, such as digital-based microlending and cryptocurrencies. The chapter contributed by Chambers, Saleuddin, and McMahan focuses on the past by examining four historical cases which, like innovations emerging today, were responding to market demand and eventually required new regulation to protect consumers.

The second section of the handbook includes five chapters examining the development of both public and private digital currencies, a topic that is of intense interest to regulators and policy-makers as this book goes to press. The chapter by Chowdry and Kim is a cryptocurrency primer providing an overview of Bitcoin and explain the mechanics of the Bitcoin blockchain with the aim of dispelling common misconceptions that have emerged regarding the decentralized process for validating Bitcoin transactions. The following chapter by Lee, Parlour, and Rajan addresses a central question in the minds of many regulators as they assess the risks and policy implications of cryptocurrencies: ‘are they economically new, or simply new technological packaging of something familiar?’ The chapter by Arjalies departs from an economic view of alternative currencies by providing a comparative analysis of three cryptocurrencies and one ‘complementary’ currency through the lens of a sociologist. Her analysis considers the production of the value attributed to each of these currencies and highlighting how their value is influenced by their social context, a reminder that adoption of technological innovation involves social processes and outcomes. The chapter contributed by Pieters addresses the ontological confusion manifest in many discussions of digital currency regulation by providing a well-structured comparison of the form and structure of digital currency types. She highlights the trade-offs that Central Banks face in deciding whether to issue a Central Bank Digital Currency (CBDC), and suggests that issuance of digital currencies by independent entities may have moved timing for CBDC outside of the control of the Central Banks. Jagtiani and her co-authors wrap up the discussion of cryptocurrencies with a chapter discussing the regulatory implications for investors, which is a timely contribution given the increasing institutionalization of private cryptocurrencies like Bitcoin and Ethereum.

The chapters in the third section of the book, Payments and Trading, are divided into two parts to provide a more coherent discussion of technological innovation on the infrastructure underpinning these exchange activities. The first part discusses alternative payments, and begins with a chapter authored by Ndung’u and Oguso that provides an empirical study of the development of digital financial services delivered via mobile phones and its impact across sub-Saharan Africa following the launch of the M-Pesa mobile payment system in Kenya. The authors are well-placed to comment on this important regional development given Ndung’u was the Governor of the Kenyan Central Bank when M-Pesa was introduced. While their chapter makes the case for regulators proactively supporting innovation they also highlight areas of emergent risk needing regulatory intervention to ensure the sustainability of public good outcomes from this innovation.

Feyen, Frost, Natarajan, and Rice continue the theme of payment innovation in developing economies by exploring the potential impact of digital money in the forms of stablecoins and central bank digital currencies in cross-border payments and remittances. The distinction between developed and developing economies is important when considering the impact of these new forms of money since the benefits they potentially bring depend on the state of the existing national payment rail infrastructure. The authors suggest that other forms of technological innovation in the payment system will potentially undermine many of the anticipated advantages advanced by proponents of stablecoins.

The second part of the Payments and Trading section explores innovation in the context of asset tokenization. Hillary and Liu begin with a review of different types of tokenized assets and highlight some of the challenges that Initial Coin Offerings introduced for regulators, including accounting, taxation and security. Gans and Gandal contribute a critical analysis of various consensus mechanisms employed to validate information recorded on distributed ledger systems, and highlight the trade-offs associated with different mechanisms. The consensus mechanisms in public blockchains are highly decentralized, which introduces challenges for regulatory supervision regimes designed for highly centralized market infrastructure. In their chapter, Matsuo, Ushida, and Wright take inspiration from internet governance models to propose a multi-stakeholder polycentric governance structure for the 'permissionless innovation' associated with decentralized systems. McKeon and Schloss survey the legal environment around blockchain applications such as smart contracts which do not fit easily within existing legal frameworks, and suggest that norms around these new applications will be the foundation for a new body of law. The application of blockchains in trade and exchange is examined in the chapter authored by Benedetti, McKeon, and Pfiffer. The authors survey the trading and exchange landscape and highlight the potential benefits that blockchain-based infrastructure can bring to these activities and identify significant barriers that need to be overcome in order to realize these benefits.

The themes of information asymmetry and coordination are addressed in the discussion of crowdfunding in Chapter Four. Chemla and Tinn begin the chapter by examining two very different forms of crowdfunding, reward-based crowdfunding and initial coin offerings, with degrees of information asymmetry. Abassi and Ivery continue the examination of information asymmetry in crowdfunding with their assessment of literature analysing the peer-to-peer online credit market, and offer their views regarding the regulatory

frameworks needed for this activity to be sustainable and contribute to a more inclusive financial system.

Schweinbacher's chapter focuses on crowdfunding and the role that online platform intermediaries play in coordinating the equity financing for early-stage companies. He compares the efficiency of these platforms with traditional sources of this funding and suggests that coordination costs remain relatively high given the uncertain returns to investors, suggesting that consolidation and collaborations with more traditional sources of equity funding may be the future direction of travel. Schweinbacher's conclusions seem prescient, given the mergers and acquisitions of crowdfunding platforms that took place during the COVID period. Chen, Kavuri, and Milne consider information asymmetry in the context of consumer protection in the regulatory frameworks for online alternative (non-bank) lending platforms in three very different regulatory jurisdictions: the UK, the US, and China. Schizas, Bennett, and Samarah round out this section with an examination of crowdfunding regulatory frameworks aimed at categorizing the approaches used by a global sample of regulatory jurisdictions. They find little consensus around 'good practice', and suggest regulatory approaches which effectively signal quality may be effective in reducing information asymmetry.

The digital transformation of the global economy is generating huge volumes of data, and the fifth section of the book examines the implications of data-driven business models emerging in financial services. Romero and Fitz begin this section by examining how psychometric analytics and artificial intelligence technologies are creating an increased level of intimacy in the relationship between providers of financial services and their customers with important implications for privacy. Romero and Fitz continue this theme in a second chapter, diving deeper into AI to provide a primer on the application of neural networks and deep learning in alternative finance, providing the reader with the opportunity to comprehend the profound implications of these technologies in the financial system.

Financial services firms based in China are increasingly recognized as leaders in big data innovation in financial services, in part because of the large population and regulatory environment providing access to the large amounts of data needed to train machine learning models. Huang and Qian examine the impact of these developments applied to credit allocation on financial inclusion and economic development in China and other developing economies. Two chapters in this section consider the impact of data-driven business models. King, Timms, and Tzameret consider how a historically

slow-moving sector, global insurance, is likely to be impacted by the adoption of big data business models, and suggest that risk of increased discrimination will be a particular challenge insurers will need to overcome given the nature of underwriting risk. Kavuri and Milne examine the data-driven business models of non-bank providers of working capital finance, essentially investigating to what extent these new platforms are substituting versus complementing traditional providers, particularly in providing access to small borrowers. This is an important issue for policy-makers deciding whether to introduce policies support innovation in alternative credit provisioning.

Data-driven investment models combining the power of smartphones with AI-enabled personalization have the potential to significantly disrupt the investment industry, and the chapter authored by D'Acunto and Rossi considers the effects of robo-advising on the investment, saving, and spending decisions made by consumers. The authors highlight several regulatory issues raised by 'democratizing' consumer access to financial advice including the risk of increasing rather than decreasing inequalities. New entrants in the investment management sector space are potential beneficiaries of open banking initiatives which enable consumers to provide third parties with access to their data held by providers of financial services. The potential of open to level the playing field between incumbents and new providers of financial services is the focus of Ziegler's concluding chapter to this section. While there is consensus around the conceptualization of open banking, Ziegler's chapter highlights how the implications for market integrity, competition, and consumer protection vary depending on the implementation regime with a comparative analysis of the UK, Australia, Hong Kong, Singapore, Brazil, and Mexico.

The sixth and final section of this handbook addresses the welfare implications of the technological innovations discussed in the prior chapters. Morse and Pence begin by examining how technology is shifting the impact of discrimination on different actors in the financial system, both positively and negatively, and highlight the critical role that policy and regulation can play in mitigating negative outcomes. Peng and Zhu discuss the implications of AI and alternative data methods used to reduce information asymmetry in trust-building mechanisms used in lending. They compare the trust-building case of China with systems in more developed countries and highlight cost-benefit trade-offs of approaches using big data analytics, particularly in countries like China with a strong trust deficit. Another factor influencing information asymmetry and trust is identity verification, which is a prerequisite for participating in the financial system. As a result, digital identity systems are a public policy priority in many countries seeking to increase financial

inclusion and mitigate risks associated with money laundering and terrorist financing. Chowdry, Goyal, and Ahmed analyse India's Aadhaar system which has provided 1.25 billion Indians with a digital identity, and offer insights regarding both the operationalization challenges and the consequences for Indian society of this system for Indian society.

The final chapter of the book authored by Boissay, Ehlers, Gambacorta, and Shin addresses an issue that may prove to be existential for many incumbent financial services firms—the entry of 'big tech' into financial services. This development also presents regulators with a number of novel challenges, particularly for those with regulatory frameworks focusing on entity-based regulation rather than activity-based regulation. The authors highlight a number of coordination issues that regulators will need to contend with as they seek to reconcile competition and privacy policy at the national level with financial regulation based on international standard-setting. The challenges introduced by the transnational nature of technological finance suggest an unprecedented need for international collaboration and coordination by national regulators.

Cambridge, UK
Cambridge, UK
Chicago, USA

Raghavendra Rau
Robert Wardrop
Luigi Zingales

Reference

Andreas Kaplan, Michael Haenlein. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25, ISSN 0007-6813. <https://doi.org/10.1016/j.bushor.2018.08.004>.

Contents

Introduction

An Overview of Technologically Enabled Finance 3

David C. Brown and Mingfeng Lin

Innovations in Alternative Finance in Historical Perspective 33

David Chambers, Rasheed Saleuddin, and Craig McMahon

Digital Currencies

A Layman's Guide to Bitcoin and Blockchain 61

Bhagwan Chowdhry and Seoyoung Kim

An Introduction to Cryptocurrencies 79

Jeongmin Lee, Christine A. Parlour, and Uday Rajan

The Role of Utopia in the Workings of Local and Cryptocurrencies 95

Diane-Laure Arjaliès

Digital Currencies and Central Banks 139

Gina Pieters

Cryptocurrencies: Regulatory Perspectives and Implications for Investors 161

*Julapa Jagtiani, Michael Papaioannou, George Tsetsekos,
Erik Dolson, and Dan Milo*

Payments and Trading

- Fintech Revolutionizing Financial Services: The Case of Virtual Savings and Credit Supply Technological Platforms in Africa** 189

Njuguna Ndung'u and Alex Oguso

- What Does Digital Money Mean for Emerging Market and Developing Economies?** 217

Erik Feyen, Jon Frost, Harish Natarajan, and Tara Rice

- Blockchain and Other Distributed Ledger Technologies in Finance** 243

Gilles Hilary and Laura Xiaolei Liu

- Consensus Mechanisms for the Blockchain** 269

Joshua S. Gans and Neil Gandal

- Multi-Stakeholder Governance in the Era of Decentralized Finance** 287

Shin'ichiro Matsuo and Ryosuke Ushida

- Law and Blockchains** 311

Stephen McKeon and Derek Edward Schloss

- Blockchain Trading and Exchange** 337

Hugo Benedetti, Stephen McKeon, and Cameron Pfiffer

- Initial Coin Offerings: Current Research and Future Directions** 369

Jiasun Li and William Mann

Crowdfunding

- How Wise are Crowds on Crowdfunding Platforms?** 397

Gilles Chemla and Katrin Tinn

- Asymmetric Information at the Platform Level: The Case of Peer-To-Peer Lending** 407

Puriya Abbassi and Raj Iyer

- Equity Crowdfunding and Capital Formation for Early-Stage Firms** 425

Armin Schwienbacher

- Growing Pains: The Changing Regulation of Alternative Lending Platforms** 441

Ding Chen, Anil Savio Kavuri, and Alistair Milne

Regulating Equity Crowdfunding: Lessons from International Practice	477
<i>Emmanuel Schizas, Thomas Kenneth Bennett, and Waleed Samarah</i>	
 Data-Driven Business Models	
The Use of Psychometrics and Artificial Intelligence in Alternative Finance	511
<i>Peter Romero and Stephen Fitz</i>	
Neural Networks and Deep Learning: A Paradigm Shift in Information Processing, Machine Learning, and Artificial Intelligence	589
<i>Stephen Fitz and Peter Romero</i>	
Use of Big Data in Developing Markets	655
<i>Yi Huang and Wenlan Qian</i>	
Use of Big Data in Insurance	669
<i>Melanie R. N. King, Paul D. Timms, and Tzameret H. Rubin</i>	
Technology and Working Capital Finance	701
<i>Anil Savio Kavuri and Alistair Milne</i>	
Robo-Advising	725
<i>Francesco D'Acunto and Alberto G. Rossi</i>	
Implementation of Open Banking Protocols Around the World	751
<i>Tania Ziegler</i>	
 Welfare Implications	
Technological Innovation and Discrimination in Household Finance	783
<i>Adair Morse and Karen Pence</i>	
Trust Building and Credit Reporting with Big Data in the Digital Age	809
<i>Wensheng Peng and Feng Zhu</i>	
Digital Identity in India	837
<i>Bhagwan Chowdhry, Amit Goyal, and Syed Anas Ahmed</i>	

Big Techs in Finance: On the New Nexus Between Data Privacy and Competition	855
<i>Frederic Boissay, Torsten Ehlers, Leonardo Gambacorta, and Hyun Song Shin</i>	
Index	877

List of Contributors

Puriya Abbassi Deutsche Bundesbank, Frankfurt am Main, Germany

Syed Anas Ahmed Indian School of Business, Hyderabad, India

Diane-Laure Arjaliès Ivey Business School, Western University, London, ON, Canada

Hugo Benedetti ESE Business School, Universidad de Los Andes, Santiago, Chile

Thomas Kenneth Bennett Cambridge Judge Business School, University of Cambridge, Cambridge, UK

Frederic Boissay Bank for International Settlements, Basel, Switzerland

David C. Brown Eller College of Management, University of Arizona, Tucson, AZ, USA

David Chambers Judge Business School, Cambridge University, Cambridge, UK;
Center for Economic and Policy Research (CEPR), Washington, DC, USA

Gilles Chemla Imperial College Business School, London, UK

Ding Chen School of Law, University of Sheffield, Sheffield, UK

Bhagwan Chowdhry Indian School of Business, Hyderabad, India

Francesco D'Acunio Carroll School of Management, Boston College, Boston, MA, USA

Erik Dolson Federal Reserve Bank of Philadelphia, Philadelphia, PA, USA

Torsten Ehlers Bank for International Settlements, Basel, Switzerland

Erik Feyen World Bank Group, Washington, DC, USA

Stephen Fitz Keio University Global Research Institute, Tokyo, Japan

Jon Frost Bank for International Settlements, Basel, Switzerland

Leonardo Gambacorta Bank for International Settlements, Basel, Switzerland

Neil Gandal Berglas School of Economics, Tel Aviv University, Tel Aviv, Israel

Joshua S. Gans Rotman School of Management, University of Toronto, Toronto, ON, Canada;
NBER, Cambridge, MA, USA

Amit Goyal Swiss Finance Institute, University of Lausanne, Lausanne, Switzerland

Gilles Hilary Georgetown University, Washington, DC, USA

Yi Huang Graduate Institute, Geneva, Switzerland;
CEPR, London, UK

Raj Iyer Imperial College London, London, UK

Julapa Jagtiani Federal Reserve Bank of Philadelphia, Philadelphia, PA, USA

Anil Savio Kavuri Loughborough University, Loughborough, UK;
Crawford School of Public Policy, Australian National University, Canberra, ACT, Australia

Seoyoung Kim Santa Clara University, Santa Clara, CA, USA

Melanie R. N. King Wolfson School of Mechanical, Electrical and Manufacturing Engineering, Loughborough University, Loughborough, UK

Jeongmin Lee Washington University, University of Arizona, St Louis, MO, USA

Jiasun Li George Mason University, Fairfax, VA, USA

Mingfeng Lin Scheller College of Business, Georgia Institute of Technology, Atlanta, GA, USA

Laura Xiaolei Liu Guanghua School of Management, Peking University, Beijing, People's Republic of China

William Mann Emory University, Atlanta, GA, USA

Shin'ichiro Matsuo Department of Computer Science, Georgetown University, Washington, DC, USA;
NTT Research Inc., Tokyo, Japan

Stephen McKeon Lundquist College of Business, University of Oregon, Eugene, OR, USA

Craig McMahon Villanova School of Business, Villanova University, Villanova, PA, USA

Alistair Milne School of Business and Economics, Loughborough University, Loughborough, UK

Dan Milo Rutgers University, New Brunswick, NJ, USA

Adair Morse Haas School of Business, University of California, Berkeley, CA, USA;
U.S. Department of the Treasury, Washington, DC, USA

Harish Natarajan World Bank Group, Washington, DC, USA

Njuguna Ndung'u African Economic Research Consortium, Nairobi, Kenya

Alex Oguso Kenya Revenue Authority, Nairobi, Kenya

Michael Papaioannou International Monetary Fund, Washington, DC, USA

Christine A. Parlour University of California, Berkeley, CA, USA

Karen Pence U.S. Department of the Treasury, Federal Reserve Board, Washington, DC, USA

Wensheng Peng Chief Economist, China International Capital Corporation Limited, Beijing, China

Cameron Pfiffer Lundquist College of Business, University of Oregon, Eugene, OR, USA

Gina Pieters Department of Economics, University of Chicago, Chicago, IL, USA;

Centre for Alternative Finance, Judge Business School, University of Cambridge, Cambridge, UK

Wenlan Qian NUS Business School, National University of Singapore, Singapore, Singapore

Uday Rajan University of California, Berkeley, CA, USA

Tara Rice Bank for International Settlements, Basel, Switzerland

Peter Romero University of Cambridge, Cambridge, UK

Alberto G. Rossi McDonough School of Business, Georgetown University, Washington, DC, USA

Tzameret H. Rubin School of Business and Economics, Loughborough University, Loughborough, UK

Rasheed Saleuddin Cambridge University, Cambridge, UK

Waleed Samarah Cambridge Judge Business School, University of Cambridge, Cambridge, UK

Emmanuel Schizas Cambridge Judge Business School, University of Cambridge, Cambridge, UK

Derek Edward Schloss Collab+Currency, Eugene, OR, USA

Armin Schwienbacher SKEMA Business School, Lille, France

Hyun Song Shin Bank for International Settlements, Basel, Switzerland

Paul D. Timms Wolfson School of Mechanical, Electrical and Manufacturing Engineering, Loughborough University, Loughborough, UK

Katrin Tinn Desautels Faculty of Management, McGill University, Montreal, Canada

George Tsetsekos Drexel University, Philadelphia, PA, USA

Ryosuke Ushida Cyber SMART, Georgetown University, Washington, DC, USA;

Fintech and Innovation Office, Japan Financial Services Agency, Tokyo, Japan

Feng Zhu Senior Research Director, Luohan Academy, Hangzhou, China

Tania Ziegler Cambridge Judge Business School, University of Cambridge, Cambridge, UK

List of Figures

A Layman’s Guide to Bitcoin and Blockchain

- Fig. 1 Worldwide Interest in “Cryptocurrency” and “Bitcoin” (Source Google Trends). In this figure, we display the weekly, worldwide Google-search interest in the terms “cryptocurrency”, “bitcoin”, and “trump” over the period spanning January 1, 2017 through January 1, 2018. The relative search interest is scaled such that 100 represents the peak popularity for a term in the given time frame for provided search-term opportunity set: {“cryptocurrency”; “bitcoin”; “trump”}. See <https://trends.google.com> for further details 62
- Fig. 2 The Bitcoin Blockchain. This figure graphically depicts a simplified version of the Bitcoin blockchain, which requires miners to find the winning “nonce” to close out a block of transactions records and to begin a new block. This nonce, when combined with the elements of the current block, must produce a double SHA-256 hash code with a minimum number of leading zeros. The first miner to find this nonce can close the current block and begin a new block, whose header will contain this nonce in addition to other pieces of information that link the new block to the prior one. Once this information is broadcast, other miners can easily verify that the new block was validly formed, and a new proof-of-work mining race begins 63

Digital Currencies and Central Banks

Fig. 1	Sweden has experienced a decline in both the value and number of banknotes and coins in circulation (<i>Source</i> The Riksbank, author calculations)	140
Fig. 2	Mobile payments and currency in circulation for Kenya are both increasing (<i>Source</i> Central Bank of Kenya, authors calculations)	141
Fig. 3	Relative Bitcoin prices move in a way that reflect the unofficial exchange rate, regardless of official exchange rate policy. This implies that individuals are using Bitcoin to circumvent capital controls and exchange rate pegs. Data Argentinian Peso to US dollar Exchange Rate: Official, Street, and Bitcoin-based. Peg in the 2013–2015 I period, moving to floating exchange rate in 2016 (<i>Source</i> LocalBitcoins pricing data, Updated authors calculations based on Pieters [2016])	155

Fintech Revolutionizing Financial Services: The Case of Virtual Savings and Credit Supply Technological Platforms in Africa

Fig. 1	Growth in micro accounts and gross deposits in Kenya (<i>Source</i> Central Bank of Kenya, 2018)	195
Fig. 2	Access to financial services and products in Kenya (<i>Source</i> FinAccess Survey 2019)	196
Fig. 3	Prevalence of savings and credit products in Kenya (% of Adult Population) (<i>Source</i> FSD Kenya 2019)	198
Fig. 4	Savings solutions in Kenya (% of Adult Population) (<i>Source</i> FSD Kenya 2019)	198
Fig. 5	Borrowing trends in Kenya (% of Adult Population) (<i>Source</i> FSD Kenya 2019)	199
Fig. 6	Growth in M-Shwari customers in Kenya (<i>Source</i> Commercial Bank of Africa, 2019)	201
Fig. 7	Growth in M-Pawa customers in Tanzania (<i>Source</i> Commercial Bank of Africa, 2019)	201
Fig. 8	Growth in MoKash customers in Uganda (<i>Source</i> Commercial Bank of Africa, 2019)	202
Fig. 9	Growth in MoKash customers in Rwanda (<i>Source</i> Commercial Bank of Africa, 2019)	202
Fig. 10	Growth in MoMoKash customers in Cote d'Ivoire (<i>Source</i> Commercial Bank of Africa, 2019)	203
Fig. 11	M-Shwari average savings and borrowings (<i>Source</i> Commercial Bank of Africa, 2019)	207

Fig. 12	M-Shwari's retained deposits, loans disbursed and non-performing loans (NPLs) (<i>Source</i> Commercial Bank of Africa, 2019)	207
Fig. 13	Age composition of M-Shwari customers as at June 2019 (<i>Source</i> Commercial Bank of Africa, 2019)	208

What Does Digital Money Mean for Emerging Market and Developing Economies?

Graph 1	Access to bank accounts and bank services is heterogeneous, but rising (<i>Source</i> World Bank Findex data; Medina and Schneider 2019)	218
Graph 2	Banking sector concentration, while declining, is associated with higher mark-ups (<i>Source</i> World Bank: Remittance prices worldwide reports)	219
Graph 3	Remittance flows are increasing (<i>Source</i> World Bank: Remittance prices worldwide reports)	220
Graph 4	Correspondent banking landscape (<i>Sources</i> Rice et al. 2020; SWIFT BI Watch; National Bank of Belgium)	221
Graph 5	Stablecoins may be supplied to a greater extent in countries with higher mobile use, internet access, remittances and trade openness (<i>Source</i> World Bank)	227
Graph 6	Stablecoins could be more attractive where cost of sending remittances is high (<i>Source</i> Rice et al. 2020 using data from SWIFT BI Watch, National Bank of Belgium; World Bank, Global Findex database and Remittance Prices Worldwide, remittanceprices.worldbank.org)	229
Graph 7	Adoption may depend on cash and card use, government effectiveness, volatility (<i>Source</i> World Bank; CPMI Red Book)	230
Graph 8	Digital payments like mobile money are already taking off (<i>Source</i> World Bank Findex data)	235
Graph 9	Digital technologies can help support inclusion and convenience (<i>Source</i> Bech et al. 2017; FIS 2018; IMF, <i>World Economic Outlook</i> , October 2019; World Bank Findex data; Instapay; national data)	236
Graph 10	FX margins make up the bulk of overall fees (<i>Source</i> World Bank)	237

Blockchain and Other Distributed Ledger Technologies in Finance

Fig. 1	Bitcoin prices	244
--------	----------------	-----

Regulating Equity Crowdfunding: Lessons from International Practice

Fig. 1	Share of regulators in Rowan et al. (2019) citing different risks among their top priorities when supervising Equity Crowdfunding—by formal remit over the sector	488
Fig. 2	Incidence of alternative finance regulation in Rowan et al. (2019)	489

The Use of Psychometrics and Artificial Intelligence in Alternative Finance

Fig. 1	SORCC Scheme of operant conditioning	535
Fig. 2	Dog classified as wolf	549
Fig. 3	Explanation	549

Neural Networks and Deep Learning: A Paradigm Shift in Information Processing, Machine Learning, and Artificial Intelligence

Fig. 1	Underfitting/high bias (left) vs. overfitting/high variance (right)	598
Fig. 2	A simple neural network	602
Fig. 3	Word2vec	608
Fig. 4	Word vector space embeddings of entire languages (French and English). The shape of the resulting structure is related to the topological properties of the underlying linguistic manifold, which is influenced by the grammar of the language. We see that French and English grammars differ significantly by observing the differences in shapes of those projections	608
Fig. 5	Neural auto-encoder	611
Fig. 6	Infinite graph representation of recurrent connection	612
Fig. 7	Long short term memory network	614
Fig. 8	Gated feedback recurrent unit	614
Fig. 9	Basic CNN diagram with max pooling	617
Fig. 10	CNN with max pooling for image processing	618
Fig. 11	Alex Krizhevsky, Ilya Sutskever, and Geoff Hinton—imagenet classification	618
Fig. 12	Filters learned by the first convolutional layer on image classification task	619
Fig. 13	Combining CNN and RNN allows relating various forms of data such as translating images into text	619

Fig. 14	Attention mechanisms can be used with both text and images to visualize the focus of the system on the pixel and character levels	620
Fig. 15	Modeling a conditional probability of the next word in a sentence using a recurrent encoder	624
Fig. 16	Transformer architecture (single time step shown)	627
Fig. 17	Generative adversarial network	628
Fig. 18	Variants of GAN architectures	629
Fig. 19	GitHub activity statistics for top Deep Learning frameworks in 2018	636
Fig. 20	Human-machine communication in natural language	638
Fig. 21	Instruction following architecture	642
Fig. 22	Long Short-Term Memory (LSTM) unit	643
Fig. 23	Virtual world dataset	646
Fig. 24	Google Street View dataset	647

Implementation of Open Banking Protocols Around the World

Fig. 1	International perspectives on Open Banking regulation and implementation	766
Fig. 2	Case-study jurisdictions by primary regulating authority and statutory objective	779

Trust Building and Credit Reporting with Big Data in the Digital Age

Fig. 1	Growing Credit Reporting Coverage (<i>Note</i> Coverage by public credit registry and private credit bureau in low- and middle-income economies versus the coverage in high-income economies. <i>Source</i> World Bank)	812
Fig. 2	Recent development of the credit reporting system in China	815
Fig. 3	Current personal credit reporting system in China	815
Fig. 4	Technological Advances Enhance Social Connectedness (<i>Notes</i> ¹ Coverage ratio is calculated as the percentage of people using mobile phones or Internet. ² Weibo, Wechat are Chinese social apps. ³ Taobao mobile is the mobile app of the largest e-commerce platform in China. ⁴ Magnetic disk price is calculated as the unit price (USD per Gigabyte) of data storage. ⁵ Internet transit prices are from drpeering.net and are US data. ⁶ Communication service price is derived from China's CPI component. <i>Sources</i> Wind, hblok.net, drpeering.net)	822

- Fig. 5 Trust is Needed When Credit Demand Grows (*Notes*
¹Annual growth rates of SME, Mortgage, and Consumption loans. SMEs are classified according to OECD standards.
²Ratios of consumption to GDP and of SME loans to total loans. Statistics of SME loans come from PBOC and SMEs are classified according to the standards of the National Bureau of Statistics. Consumers covered by the Credit Reference Center at the People's Bank of China (CCRC) is calculated as the percent of total population. ³Scaled by the total population. ⁴Scaled by the number of market entities published by State Administration for Market Regulation. *Sources* OECD; PBOC; Wind) 825
- Fig. 6 E-commerce, Digital Payment and Digital Financing (*Notes*
¹Ratio of online sales as % of retail sales of consumer goods; total volume of third-party payment. ²WeBank and MYBANK new loans. ³Ratio of P2P trading volume to new bank loans. *Source* Analysys; PBOC; WDZJ.com; Wind) 826

Big Techs in Finance: On the New Nexus Between Data Privacy and Competition

- Graph 1 Big techs' life cycle: Theory and practice. ¹The firm's life cycle described in the left-hand panel borrows from the synthesis of the literature by Miller and Friesen (1984). Given that big techs are still new and rising firms, we purposely ignore the usual "decline" phase. MSP = multi-sided platform (*Source* BIS [2019]; Miller and Friesen [1984]; S&P Capital IQ; Authors' calculations) 858
- Graph 2 Big tech credit, asset prices and bank development. Robust standard error in parentheses. **indicates significance at the 5% level. ¹The ratio is calculated for 2017 and is defined as big tech credit divided by total credit to the private non-financial sector (including total fintech credit). ²Average over the period 2010–2015. ³Period of estimation: 2005–2013. ⁴Period of estimation: 2011–2017 (*Source* BIS [2019]; Frost et al. [2019]; Gambacorta et al. [2019]; World Bank; Cambridge Centre for Alternative Finance and research partners; company reports; Authors' calculations) 860

- Graph 3 Credit assessment and big data analytics. ¹The loss rate is the volume of loans more than 30 days past due relative to the origination volume. To date, Mercado Libre's internal rating system has proved better able to predict such losses. It segments loan originations into five different risk groups as compared with the three clusters identified by the bank bureau. The size of the dots is proportional to the share of the firms in the rating distribution. ²True positive rates versus false positive rates for borrowers at different thresholds for a logistic model with only the credit bureau score (I), a logistic model with the bureau score and borrowers' characteristics (II), and a machine learning model with the Mercado Libre credit score (III). A random model is included for comparison purposes. The ROC curve shows that the machine learning model has superior predictive power to both the credit bureau score only and the credit bureau score with borrower characteristics (*Source* BIS [2019]; Frost et al. [2019]) 863
- Graph 4 Fintech credit scores show greater predictive power after a change in regulation. The vertical dashed line indicates when the People's Bank of China (PBoC) issued specific draft guidelines to tighten regulations on shadow banking. In particular, from 17 November 2017, financial institutions have not been allowed to use asset management products to invest in commercial banks' credit assets or provide "funding services" for other institutions (such as fintech companies) to bypass regulations. The new rule has had a huge impact on fintech companies' funding sources. The PBoC set also a limit on the interest rates charged by P2P lending companies. All annualised interest rates, which include the upfront fees charged for loans, were capped at 36%. The effects of these new rules were also reinforced by the strict measures concerning online micro-lending that were imposed on December 1, 2017 by China's Internet Financial Risk Special Rectification Work Leadership Team Office. ¹The vertical axis reports the Area Under the ROC curve (AUROC) for every model. The AUROC is a widely used metric for judging the discriminatory power of credit scores. The AUROC ranges from 50% (purely random prediction) to 100% (perfect prediction) (*Source* Gambacorta et al. [2019]) 864
- Graph 5 From market competition to market manipulation (*Source* Adapted from Bar-Gill [2019]) 867

- Graph 6 Willingness of users to share data for better offers in financial services varies. ¹The question in the survey reads “I would be comfortable with my main bank securely sharing my financial data with other organisations if it meant that I received better offers from other financial intermediaries”; for Belgium the figure is calculated over Belgium and Luxembourg (*Source* EY [2019]; IMF, *World Economic Outlook*; Chen et al. [2020])

List of Tables

Innovations in Alternative Finance in Historical Perspective

Table 1	Number of IPOs on the LSE since 1900	43
Table 2	IPO failure rates	45

An Introduction to Cryptocurrencies

Table 1	Daily (top) and monthly (bottom) altcoin correlations with other assets	85
---------	-------------------------------------------------------------------------	----

The Role of Utopia in the Workings of Local and Cryptocurrencies

Table 1	The valuation infrastructure of Sol Violette, Bitcoin, Ğ1 “June” and Impak Coin	105
Table 2	The three functions of money of Sol Violette, Bitcoin, Ğ1 “June” and Impak Coin	128

Fintech Revolutionizing Financial Services: The Case of Virtual Savings and Credit Supply Technological Platforms in Africa

Table 1	Characteristics of the virtual savings and credit products	205
Table 2	Some regulated digital lenders and unregulated app-based fintech lenders in Kenya as at October 2018	211

What Does Digital Money Mean for Emerging Market and Developing Economies?

Table 1	Categorisation of cash, crypto-assets, stablecoins, bank deposits and CBDCs	223
Table 2	Factors that may support the adoption of digital money	226
Table 3	Particular challenges of stablecoins for EMDEs	231

Blockchain and Other Distributed Ledger Technologies in Finance

Table 1	Strengths and weaknesses of blockchain technology	244
---------	---------------------------------------------------	-----

Initial Coin Offerings: Current Research and Future Directions

Table 1	Below lists several empirical papers and the datasets they use	379
---------	----------------------------------------------------------------	-----

Growing Pains: The Changing Regulation of Alternative Lending Platforms

Table 1	Lending volumes: consumer marketplace lending	442
Table 2	Lending volumes business marketplace lending	443

Regulating Equity Crowdfunding: Lessons from International Practice

Table 1	Risks to equity crowdfunding Investors	486
Table 2	Regulatory requirements in regulated equity crowdfunding markets, by model of regulation (Rowan et al 2019)	493
Table 3	Mitigation strategies for key risks in equity crowdfunding, by regulatory model	498
Table 4	Summary statistics on regulatory requirements by regulatory model	500

The Use of Psychometrics and Artificial Intelligence in Alternative Finance

Table 1	Selection of ambient data with sources and use cases	569
---------	------------------------------------------------------	-----

Use of Big Data in Insurance

Table 1	Five characteristics of big data in the context of insurance	673
---------	--------------------------------------------------------------	-----

Technology and Working Capital Finance

Table 1	Alternative finance total and SME lending in the most important jurisdictions	707
Table 2	The variety of working capital finance	710

Introduction



An Overview of Technologically Enabled Finance

David C. Brown and Mingfeng Lin

1 Introduction

Information is fundamental to all businesses. Generating information about consumer demands creates market opportunities and information about industry dynamics leads to competitive advantages. Processing information leads to improved decision making, creating efficiencies, and reducing costs. Transmitting information allows individuals and businesses to communicate and interact, facilitating valuable trade. Recent advancements in information and internet technologies have changed how we generate, process, and transmit information, and in doing so, have changed how businesses operate. Within the finance industry, the use of these technologies to create new financial services or to improve existing services, is broadly referred to as technologically enabled finance, or “FinTech.”

Harnessing the value of information faces both technical and economic challenges. From a technical standpoint, information may be hard to process and incorporate into decision processes, like “soft information” in the

D. C. Brown (✉)

Eller College of Management, University of Arizona, Tucson, AZ, USA

e-mail: dcbrown@email.arizona.edu

M. Lin

Scheller College of Business, Georgia Institute of Technology, Atlanta, GA, USA

e-mail: mingfeng.lin@scheller.gatech.edu

banking industry (Liberti and Petersen 2018). From an economic standpoint, owners/producers of information may not be willing to give away information due to privacy or financial concerns, and receivers of information need to have reasons to trust its contents. New FinTech innovations can help to overcome both challenges, by changing the ways information is generated, processed, and transmitted through our financial system.

As an example, a typical argument for the existence of community banks is their ability to gather and process soft information about local borrowers and businesses. By interacting with borrowers directly prior to and after loan origination, community banks are able to gauge their creditworthiness beyond the typical credit information vector. Information technologies such as artificial intelligence, particularly text mining, computer vision, and video mining, have the potential of significantly expanding the number of quantitative dimensions upon which a borrower can be evaluated. Transferring information about a particular borrower is also much cheaper than person-to-person communications based on personal perceptions and subjective judgment. Ignoring privacy and security concerns (for now), blockchain technology can take this one step further to make errors about a borrower's credit profile much less likely, further reducing the cost of transmitting information (from one bank to another, for example). In fact, by ensuring the accuracy of information through technologies such as blockchain, technology can do increase the "trust" that users have about information. Increasing trust can mitigate information asymmetry when information is transferred from the original stakeholders (e.g., the original borrower and lender) to future stakeholders who rely on historical information (e.g., information in credit reports) for future decision-making.

Some of the most exciting changes that gave rise to the FinTech phenomenon have roots in how technologies enable gathering information from large networks of people. Crowdfunding is a prime example. Even if we look at the "simplest" type of crowdfunding—donation-based crowdfunding such as Kiva.org (it is relatively simple since there are fewer financial implications than other types)—it would have been technically impossible and financially prohibitive to reach potential donors on a global scale and fund small loans in the magnitude of just several hundred or thousand dollars each. Thanks to online platforms, it becomes much more likely for small-scale donations to create significant social impacts by matching the demand and supply sides of funds. Rewards-based crowdfunding goes a step further allowing enthusiasts to provide funding for early stage business ideas through pre-sales of products not yet produced. Debt-based and equity-based crowdfunding—more squarely in the realm of finance—match the supply and

demand of capital in expectation of financial returns. Crowdfunding platforms enable transactions that were unthinkable before Web 2.0 (and beyond) allowed interactive and social content (compared to the static internet in Web 1.0), and new technologies will continue to be used to bring together the supply and demand sides of funds.

From the perspective of personal finance, tremendous opportunities have been identified, and many remain to be explored, in connecting individual consumers to the “right” mix of financial assets. Robo-advisors have become prominent, and both financial planning and insurance are embracing new technologies. In fact, the scope of “financial assets” has significantly changed in recent years, particularly fueled by the astonishing rise—and subsequent volatility—of cryptocurrencies. Determining how to incorporate such new assets classes into individuals’ financial management presents significant opportunities for researchers and businesses alike.

Despite tremendous growth thus far, FinTech has yet to realize its full potential. Most notably, tensions remain between the proclaimed benefits of technologies, and various concerns of using those technologies. For example, FinTech often utilizes “big data,” but big data presents inherent privacy concerns. Balancing the competing goals of enhanced information quality and tight privacy controls creates regulatory challenges and new business opportunities.

In this chapter, we focus on how financial technologies have changed the way that market-relevant information is gathered, produced, verified, stored, transmitted, and processed. We consider how financial technologies enable a more efficient marketplace, and also how efficiency gains may raise privacy and regulatory concerns. We organize our discussions in two main areas. The first section focuses on the matching between supply and demand of capital, including individual credit access and business financing. We discuss new business models such as peer-to-peer lending and other types of crowdfunding, and the impact of new technologies such as blockchain on record keeping and credit monitoring. The second section examines FinTech from the perspective of personal finance, discussing new investment opportunities enabled by the growth of FinTech and new ways business are providing financial guidance to investors. Throughout the chapter, we highlight new businesses that are driving FinTech innovations and existing businesses that are adopting new technologies to remain competitive. We also discuss regulatory and social welfare issues, examining the trade-offs between privacy and efficiency gains as well as the potential for biases in algorithms. Finally, we discuss unresolved questions that provide potentially fruitful research opportunities.

2 Matching Supply and Demand of Capital

2.1 Individual Credit Access

The growth of data, the Internet, and computational machinery have significantly affected how individual creditworthiness is quantified. We now know much more about individual consumers. In addition to the hard, cold numbers in credit reports, companies have collectively accumulated a potential gold mine of data about nearly everyone. From social media activities (Facebook, Twitter, Pinterest), to the physical trajectory of mobile phones and every single touch on those phones; from order histories on e-commerce websites, to health expenditures, the amount of information that businesses (individually and collectively) gather about consumers is staggering. In this section, we explore ways this transformation has affected individual credit access, and what questions these effects raise for both regulators and researchers.

A consumer's online friendship network, typically revealed in his or her social networks, can serve as a useful mechanism to gauge his or her creditworthiness. Online person-to-person lending was one of the earliest industries to explore the power of social networks in enhancing credit access, either in the sense of reducing adverse selection or reducing moral hazard. [Prosper.com](#), one of the earliest such platforms in the United States, allowed users to create their friendship networks through a referral process, and research shows that lenders can use network features to infer borrowers' quality (Lin et al. 2013). [LendingClub.com](#) started as an app on Facebook, and even after moving onto its own website, its platform promoted social networks among borrowers and lenders for years. The platform provided a tool for lenders to identify potential connections to the borrower (location, employer, alma mater, etc.). Other aspects of information revealed by the borrowers can also help predict their willingness and ability to repay. For example, linguistic features of text written by borrowers can be "mined" to predict borrowers' default likelihood (e.g., Gao et al. 2020).

Some international FinTech startups have pushed the envelope even further. In China, for example, a growing number of companies originate small loans to consumers based on a virtual "dump" of personal information on their mobile phones. WeCash.net and Qudian are two such examples. The big data obtained from these borrowers can include intimate personal data such as e-commerce activities and browsing histories, bank account information, cell phone location trajectory, and even a borrower's list of

contacts—your information could be in their database if your friend requests funds from them.

As algorithms improve, we will increasingly be able to extract economically meaningful information from vast amounts of structured and unstructured data about potential borrowers (text, images, videos, etc.), and even their online friends. But what do these changes imply for society?

From an economic point of view, how big data will affect lenders (both FinTech startups and incumbent banks) and borrowers is unclear. Increased data access could be great news for lenders, because they can now get closer to first-degree price discrimination. With access to more accurate information about borrowers, lenders can price loans to potentially reach borrower's maximum willingness to pay. From borrowers' point of view, better information and the potential for enhanced price discrimination will likely lead to benefits and costs, and create both winners and losers in the process. Worse borrowers who were able to "blend in" with better borrowers in the past will now either have to pay higher interest rates, or may not be able to borrow at all. Separating from worse borrowers (whether through characteristics in a credit report or behaviors such as driving patterns) will likely improve better borrowers' rates, but how much those borrowers benefit will depend on the degree of lender competition. Increased data availability may reduce lenders' costs and enhance competition for borrowers' benefit (e.g., by attracting more lenders to enter the market), and it may also enable some credit-worthy borrowers who did not have access before to gain access due to the alternative data sources made possible through big data. Furthermore, if borrowers can easily share their information with lenders, rather than lenders having control of their information, borrowers may be net beneficiaries of big data through less "lock-in" and increased competition. In either case, the platforms and providers of technologies will certainly try to obtain a share of the surplus gain from reduced information asymmetry; but how this will affect traditional capital providers (such as banks) is less clear. FinTech startups often work closely with these traditional institutions; for example, LendingClub, [Prosper.com](https://www.prosper.com) and Kabbage all collaborate with banks in the issuance of debt notes. Such relationships are likely to benefit some banks, but the effect on the "average" traditional bank is ambiguous.

From a regulatory standpoint, it is even less clear how big data will, and should, be used. To start, algorithms are often either proprietary to the platform or lender, or licensed through third parties, making monitoring the appropriateness and legal compliance of such algorithms challenging. For example, will regulators be able to ensure that lenders are not using

protected information such as gender, age, and race in making credit decisions?¹ One potential means to monitor compliance is by evaluating a lender's loan portfolio; do the borrowers' characteristics match the populations' characteristics? While lenders may not explicitly use protected information, it seems inevitable that at least some of the big data about borrowers could significantly correlate with those protected characteristics. In that case, does using information that is correlated with protected information, such as a borrower's neighborhood or zip code, constitute discrimination? Alternatively, if a lender's portfolio characteristics do not naturally match population characteristics, should the lender be able to use protected information (or information correlated with protected information) to target particular customers such that their resulting portfolios do match population characteristics? How firms and regulators should and will address these issues are challenging and meaningful questions to pursue.

2.1.1 Microlending

Microlending typically refers to donation-based crowdfunding for entrepreneurs in third-world countries. A well-known platform is Kiva.org, which enables donors to help entrepreneurs in impoverished countries.² The persons providing funding are considered "donors" because, although the borrowers still have an obligation to repay, the interest rate that the donors receive is zero. Kiva makes this possible by working with microfinance institutions (MFIs) around the globe, who work with local entrepreneurs, obtain information from them, and post their requests online. Once the funds are sent to the entrepreneurs, the microfinance institutions monitor the repayment process and provide updates to the donors.

Microlending has some significant differences from donations, however. Very importantly, although donors do not charge interest rates on their donations, the microfinance institutions do. According to Kiva, the microfinance institutions charge an average interest rate of 35%, which often does not cover the high operating costs of these lenders.³

Microlending demonstrates the power of internet technologies in matching the supply and demand of funds on a global scale. But questions remain about the actual economic impact of such activities, especially

¹ Butler et al. (2020) shows evidence of racial discrimination in the auto loan market.

² Recently, Kiva has also started working with entrepreneurs in the US.

³ <https://nextbillion.net/kiva-responds/>.

the distribution of gains from this process. For example, many microfinance institutions are profit-driven, and the monitoring of their activities remains challenging due to the varying legal structures and rights in different countries (Yermack 2018). This is exacerbated by the fact that the ultimate borrowers (entrepreneurs in economically disadvantaged areas) may not have internet access without the help from the MFIs. How should these MFIs be regulated? When is their profit excessive? To what extent should donors believe the information that they see on internet platforms, from the descriptions of borrowers, to the status updates? Even if information is fabricated, it is virtually impossible for other parties to verify. None of these questions are easy to answer, as demonstrated by stories of MFIs on Kiva abusing the goodwill of donors and committing fraud.⁴ If left unregulated, or under-researched, microlending may not realize its potential to address the needs of those at the bottom of the pyramid. Microlending could even strengthen the market power of potential loan sharks by providing them with cost-free capital, further hurting disadvantaged demographics.

2.1.2 Peer-to-Peer Lending

Peer-to-Peer (P2P) lending⁵ originated as an almost idealistic model of finance: circumventing big banks and corporations (the so-called “middle man”) so that borrowers can pay less and small investors can earn more.⁶ P2P lending is one of the first major FinTech innovations, and in terms of scale, it has always been the largest among all types of crowdfunding.⁷ Institutional investors now play an increasingly large role in P2P markets, so the word “peer” may soon lose its relevance (Lin et al. 2018; Vallee and Zeng 2019). Moreover, there is significant concern over whether the industry’s growth is sustainable. For example, the UK FCA (Financial Conduct Authority) is revisiting whether retail investors are able to successfully take advantage of this new market and are investigating whether retail investors were sufficiently warned about the underlying risks in this industry.⁸

⁴ http://media.kiva.org/INNOV-SKOLL-2009_flannery.pdf.

⁵ There is a small but growing academic literature on peer-to-peer lending. Examples include Chava et al. (2017), Pope and Sydnor (2011), Freedman and Jin (2008), Hildebrand et al. (2016), Tang (2019), Duarte et al. (2012), Iyer et al. (2015), Lin and Viswanathan (2015), Wei and Lin (2016), Michels (2012), and Zhang and Liu (2012).

⁶ While P2P lending may avoid banks, Tang (2019) shows that credit expansion from P2P lending likely occurs only among borrowers who already have access to bank credit.

⁷ <http://crowdexpert.com/crowdfunding-industry-statistics/>.

⁸ <https://www.ft.com/content/8333421e-9186-11e8-b639-7680cedcc421>.

From a regulatory point of view, monitoring can be especially challenging when platforms have a peer-to-peer component. For example, whether an individual is able to borrow funds from a platform, and the interest rate at which he or she borrows at, are not solely determined by the platform or the algorithm that it employs. After a proprietary algorithm determines the creditworthiness of the individual, lenders (retail investors and institutional investors) still decide whether or not to lend to a particular individual. So even if some individuals are (inappropriately) disadvantaged due to their protected status, it is virtually impossible to fully attribute that to a platform's algorithm. In the context of [Prosper.com](#), research has shown that black borrowers were less likely to secure funding from [Prosper.com](#) (Pope and Sydnor 2011) and beautiful borrowers pay a lower rate (Ravina 2019). However, it is virtually impossible to attribute that to an algorithmic bias; it could very well be the bias on the part of investors, not [Prosper.com](#) itself. In such cases, should the platform be liable? How can researchers test whether discrimination is due to a platform's bias or investors' biases?

The history of P2P lending in the Chinese market raises other interesting questions. In the Chinese market, P2P lending witnessed dramatic growth, and a swift decline as regulators stepped in. Due to the size of the Chinese market, P2P lending grew at an astonishing rate prior to 2017, reaching as many as 3500 platforms in 2015.⁹ The intense competition among platforms led many of them to resort to all types of “novel” mechanisms to recruit investors and borrowers. One such mechanism was the “promise” that investors' principal is safe: for a period of time, many platforms explicitly or implicitly advertised that investors' principal was guaranteed. As a result, as much as 40% of these platforms were basically Ponzi schemes¹⁰—using funds from new lenders as “returns” for new investors—which created a bubble waiting to burst, and chain reactions among some platforms that were financially related to each other. After several high-profile P2P lending platforms went bankrupt, regulators clamped down, and many more platforms went out of business.¹¹

The experience in the Chinese market raises several questions that are relevant to the global FinTech industry. Is a “light touch” approach of regulation—letting platforms grow and dealing with problems later—advisable? Would the opposite approach have killed such innovations? What was the turning point that transformed an idealistic notion into a Ponzi scheme?

⁹ <https://www.finextra.com/blogposting/17107/the-rise-and-fall-of-p2p-lending-in-china>.

¹⁰ <http://tfageeks.com/2018/08/20/p2p-lending-crisis-china-will-end-p2p-lending/>.

¹¹ <https://www.ft.com/content/75e75628-8b27-11e8-bf9e-8771d5404543>.

What kind of research or regulations can help avoid such dramatic cycles to play out in other innovations?

2.1.3 Soft Information and Banks

FinTech has the potential to “harden” soft information (Stein 2002; Liberti and Petersen 2018), moving soft information into the realm of big data and algorithmic decision making. Community banks have traditionally had an advantage in processing soft information, allowing them to establish niche markets. This has been particularly important in lending to low-income households (Mayer 2018). But such advantages may be weakened due to the rise of FinTech, as competition may come from larger banks as well as FinTech startups that specialize in turning soft, unstructured information into hard, quantifiable information.¹² Big banks may also have an advantage in leveraging FinTech due to their scale, the size of their customer base (a requirement of big data), and their financial resources.¹³ Once new technologies are established and proven, large banks could ultimately leverage those technologies as well, thereby chipping away at the soft-information-based competitive advantage of community banks.

As an example, startups like [Kabbage.com](https://www.kabbage.com) are making small business loans using non-conventional data about supply chains, accounts payable, and accounts receivable. The high valuations of similar startups is some evidence that this is a promising venue of business growth. It is not yet clear how community banks will respond to these FinTech trends, and whether they will discover new ways to differentiate from large banks—or perhaps collaborate with FinTech startups. By hardening soft information, will large banks better serve low-income households and other under-banked populations? This is an area that empirical research may be able to detect a trend, and it could have significant implications for the competitive landscape of the banking industry, which in turn will affect consumer and social welfare.

¹² Fuster et al. (2019) shows that FinTech lenders process mortgage applications faster without incurring more defaults, but finds no evidence that those lenders target borrower with low access to finance.

¹³ There is also some evidence that the introduction of big banks drives out soft information and reduces lending for individuals and small businesses (Rice and Strahan 2010; Cornaggia et al. 2015).

2.1.4 Record Keeping

Traditional financial transactions often suffer from inaccuracies due to record-keeping errors. For example, a number of homeowners have had their houses sold or foreclosed on by banks that had no claim to the property.^{14,15} Blockchain technologies, by being decentralized and secure, have the potential to eliminate such errors. But there is a long way to go before that promise is realized.¹⁶ To realize the potential of blockchain technologies, government agencies (local, state, and even federal), mortgage companies, real estate agents, and the judicial system must recognize and require that the record maintained through blockchain is the only valid proof of ownership. Once that goal is reached, such errors will be minimized, but the transition from the current failed, patchwork system to a blockchain-based system will take a long time. Technical complexities need to be solved (including determining which blockchain technology will emerge as dominant) and consideration should be given to the effects on current stakeholders in record-keeping processes. Lawyers, title companies, and staff members in government entities that currently maintain such records will likely resist change, further slowing adoption of blockchain technologies.

2.1.5 Credit Monitoring

Another apparent venue where blockchain or other decentralized ledger technologies could be extremely valuable is personal credit information. Currently, credit reporting agencies have a legislative mandate to collect and store vast amounts of personal credit information. Those agencies benefit from their near-monopolistic market power by selling that information to financial institutions and other businesses. Such power also makes them subject to intense cybersecurity attacks because they are information “hubs.” As long as these agencies are the central repository for personal credit information, Equifax will not be the last agency to deal with a data security breach.¹⁷ Further, even if data security were not an issue, the credit agencies often harm consumers’ abilities to access credit, even if inadvertently, by reporting outdated or inaccurate information.¹⁸ Using blockchain-based

¹⁴ <https://www.ft.com/content/b195fb02-2fde-11e7-9555-23ef563ecf9a>.

¹⁵ <https://abcnews.go.com/Business/bank-america-sued-foreclosing-wrong-homes/story?id=9637897>.

¹⁶ Yermack (2019) discusses the current state of blockchain technology in finance.

¹⁷ <https://www.consumer.ftc.gov/blog/2019/07/equifax-data-breach-settlement-what-you-should-know>.

¹⁸ <https://www.cnn.com/2017/09/27/the-real-problem-with-credit-reports-is-the-astounding-number-of-errors-equifax-commentary.html>.

technologies to store personal credit data has the potential to make cyber attacks much more difficult, and by giving individual consumers easier access their credit information, the potential to reduce the likelihood of reporting errors.

While storing personal credit information on blockchain-based platforms has potential benefits, there may also be downsides. One justifiable benefit of the traditional reporting agencies is their roles as “gatekeepers.” To some extent, the agencies protect consumers’ privacy by determining who (other than the consumers themselves) has access to consumers’ credit information. New technologies based on blockchain will have to account for the delicate balance between decentralization and privacy.

2.2 Payment Systems

Another significant focus of FinTech startups is the transfer of funds, particularly person-to-person transfers. In the past, consumers primarily used cash, bank transfers, checks, credit cards, or other payment systems to pay other consumers or businesses. The growth of information technologies, especially the wide diffusion of personal smart phones, has led to significant growth in this new FinTech area.

The most prominent example of person-to-person monetary transfers is perhaps WeChat Pay, which has transformed the way consumers interact with businesses and each other in China. While WeChat started out as a social networking app, it eventually allowed users to link their WeChat account to their bank accounts. In doing so, consumers could make payments to other individuals or companies (such as grocery stores) by scanning payment codes (QR codes) or through in-app purchases. The adoption of WeChat pay (and to a lesser extent AliPay, a payment system offered by the Alibaba Group) has been so dramatic that China has become an (almost) cashless society—even beggars can accept WeChat payments.¹⁹

In the United States, adoption of new payments technologies has been slower and competition has fragmented the market. Most prominently, Venmo provides a similar service, but its penetration is nowhere near that of WeChat in China. Facebook, Apple, and Google have all started similar services to facilitate interpersonal transfers, and even banks started offering similar services such as Zelle. Such intense competition and the resulting fragmentation has further diminished the usefulness of these methods. WeChat and AliPay were able to establish an undisputed duopoly for payment

¹⁹ <https://www.mirror.co.uk/tech/beggars-now-accepting-mobile-payments-11700364>.

methods in a relatively short period of time—one beneficial aspect of such a market structure is that almost everyone accepts WeChat and/or AliPay as a payment method in China.

The adoption of cashless payment systems has significant implications for regulators and researchers. Through WeChat, for example, a private company (and potentially the government) has a much better understanding of the shadow economy, which has always been very difficult to measure and monitor. Will cashless transactions become monitored and taxed, potentially leading consumers toward more opaque cash transactions? How will cashless payment systems affect monetary policy, and how does that relate to the competitive landscape of the payment platforms?²⁰ If payment systems lead to a natural monopoly, should it be regulated in the same ways as traditional natural monopolies, such as utilities companies? What is the role of traditional banks in this transformative process?

2.3 New Venture Financing

Successfully financing new ventures depends on overcoming several obstacles between potential investors and entrepreneurs (firms). First, there is often separation between potential investors and entrepreneurs—either physical or virtual—that prevents entrepreneurs from finding the “right” investors. Traditionally, entrepreneurs have been limited to soliciting capital from nearby investors (giving rise to large start-up hubs) or utilizing matchmaking services to provide access to investors. Second, there is often an information gap between entrepreneurs and investors. Entrepreneurs often possess more information about their internal operations, experience, and product development efforts. Based on this information, entrepreneurs ultimately have a better (even if still imperfect) idea of the quality of their firm, and the likelihood of the firm succeeding in the future. This traditional information asymmetry problem is often resolved through the choice of financing raised (debt vs. equity), evaluation by professional investors, or certification by financial intermediaries. FinTech is being used to overcome these obstacles by better matching entrepreneurs and investors and reducing frictions due to information asymmetry.

²⁰ <https://nextbillion.net/mobile-money/>.

2.3.1 Crowdfunding

Similar to microlending and peer-to-peer lending, crowdfunding uses the Internet to raise money from a large number of people for a common purpose. While crowdfunding can be used to fund charitable causes or artistic projects, in a business context the purpose is often financing a new venture (i.e., a start-up company), either through pre-selling products or services, or through selling securities that give cash flow and/or control rights to investors. By linking entrepreneurs to a new pool of investors, crowdfunding can harness the “wisdom of the crowd.”

From a finance perspective, the most traditional use of crowdfunding involves sites that cater to sophisticated investors. AngelList and SeedInvest both provide online venues for sophisticated investors to be paired with investment opportunities. As sophisticated investors, these individuals are thought to be more capable of analyzing companies and assessing the quality of management and companies’ potential for success. Either the platforms themselves or large lead investors such as venture capitalists serve as certifiers of the entrepreneurs, effectively curating a selection of investment opportunities for investors. In these cases, the Internet serves as an efficient means of bringing together investors and entrepreneurs, but the parties play very similar roles as would be found in investment banking or venture capital, just at a smaller scale.

The defining characteristic of most crowdfunding sites is that anyone can have access to the entrepreneurs and investment opportunities. Expanding the pool of investors brings many opportunities, but is not without risks. Sophisticated individuals may have the knowledge and time to invest in potential deals, but less-sophisticated investors who contribute relatively small amounts cannot be expected to do the same. As a potential means to protect such investors, almost all crowdfunding platforms have adopted all-or-nothing financing rules. Platforms require that entrepreneurs set financing thresholds, and if pledged capital does not meet or exceed the thresholds, all capital is returned to investors. Ideally, all-or-nothing rules help to protect investors and aggregate the opinions of a large, diverse group of investors. As we will discuss, the efficacy of all-or-nothing rules depends on the types of crowdfunding campaigns and investors involved.

Before discussing different crowdfunding types in more detail, it is important to highlight that all crowdfunding platforms strive to provide efficient financing. By efficient financing, we mean that positive-valued projects exceed their financing thresholds and proceed, while negative-valued projects do not receive sufficient funding and are canceled. In this way, crowdfunding

platforms and their investors act as a screening mechanism, only providing capital to worthy projects. The more efficient the platform, the better for entrepreneurs (who are more likely to get capital for good ideas, and less likely to waste time pursuing bad ideas) and the better for investors (who are less likely to fund bad investments). As a result, many questions surrounding crowdfunding relate to how platforms can be as efficient as possible.

One way to improve financing efficiency is to reduce asymmetric information between entrepreneurs and investors. New ventures' traditional asymmetric information problem is that entrepreneurs have better information about the firm's technology, operations, and prospects. As with AngelList and SeedInvest, many platforms address this problem by acting as certifiers and monitoring the progress of companies. The biggest challenge relative to traditional financing is that the deals are small, and so are investors' stakes in the firms. Holding only small stakes, individuals have little incentive to monitor firms or protect their own interests in future financing rounds. Many platforms address this issue by holding relatively large stakes in the firm or warrants that provide substantial upside. The potential for large gains provides greater incentives for monitoring, which benefits small investors.

New ventures also face an additional source of asymmetric information when using crowdfunding. In many cases, the crowd of investors may have incremental information that is critical to the business's success. For example, if the crowd of investors overlaps with the business's potential customers, the fundraising process can be a strong indicator of product demand. A product with strong demand will likely solicit contributions in excess of its all-or-nothing threshold, while less popular products may fail to secure adequate funding. By learning from the wisdom of the crowd, entrepreneurs, investors, and crowdfunding platforms all benefit.

2.3.2 Rewards-Based Crowdfunding

Rewards-based crowdfunding, popularized by Kickstarter and Indiegogo, operates by entrepreneurs' pledging rewards to consumer-investors based on their level of contribution. In many cases, rewards-based crowdfunding can be thought of as a product pre-sale. Compared to traditional financing, it is unique in that the investors are also the consumers of the product. Because consumers finance the project, fundraising totals reflect the demand for the product, providing a strong signal to entrepreneurs. Importantly, by eliciting the wisdom of the crowd, product pre-sales can be a very efficient form of financing. Theoretically, Ellman and Hurkens (2019) and Strausz (2017) show that an optimal crowdfunding mechanism can successfully aggregate

the wisdom of the crowd when capital comes from consumer-investors who have private values for the products they purchase in pre-sales. Schwienbacher (2018) and Chemla and Tinn (2020) also analyzes models in which reward-based crowdfunding successfully aggregates the wisdom of the crowd.

Many of the projects funded via rewards-based crowdfunding have produced consumer products such as electronic watches, computer hardware, and video and board games. As consumer products are only part of the economy, it is natural to ask whether rewards-based crowdfunding is limited to such products or whether it can be applied to business products or services. What product characteristics are important for successful rewards-based crowdfunding? Can small, local businesses, such as restaurants or breweries, successfully finance themselves via rewards-based crowdfunding?

While rewards-based crowdfunding appears well-suited to aggregating information from consumer-investors, it is less clear how effective platforms are at mitigating the more traditional forms of information asymmetry. Enforcing product quality may be particularly hard when large quantities of product are pre-sold, creating pressure to fulfill the orders quickly. Entrepreneurs may also be tempted to deliver to new customers before fulfilling pre-sold orders, and some may even renege on their promises leading to fraudulent campaigns.²¹ Whether these issues become apparent and how they may be mitigated are potential directions for future research.

2.3.3 Securities-Based Crowdfunding

When a new venture's product does not lend itself to pre-sales, firms can turn to security-based crowdfunding, in which they sell claims to the firms' cash flows and control rights in exchange for capital. Securities-based crowdfunding is less developed than rewards-based crowdfunding, at least partially due to the regulatory requirements surrounding the sale of securities. While securities-based crowdfunding has been active in the UK and Europe since 2013, US securities-based crowdfunding has only been legal since 2016. Given the later start, and the longer horizon required to measure success, data, and research on securities-based crowdfunding is rather limited.

A theoretical literature on securities-based crowdfunding is developing, as several papers have focused on the financial incentives of investors who will not necessarily consume a venture's product. Brown and Davies (2020) shows that when investors are strictly profit-motivated and invest together,

²¹ Some campaigns have turned out to be outright fraudulent, such as a board game campaign on Kickstarter. See <https://www.ftc.gov/news-events/press-releases/2015/06/crowdfunding-project-creator-settles-ftc-charges-deception>.

securities-based crowdfunding cannot efficiently aggregate the wisdom of the crowd. Rather, as crowds become larger, financing becomes less efficient. The inefficiency arises as feedback from the wisdom of the crowd. If the entrepreneur can use the information in fundraising amounts to cancel bad projects, then the investors anticipate bad projects will be canceled and contribute regardless of their information. Fundraising amounts then become uninformative, and financing is inefficient. In contrast, Cong and Xiao (2019) shows that securities-based crowdfunding can be efficient due to the existence of all-or-nothing financing thresholds. When investors arrive sequentially, they learn from the actions of all prior investors and the all-or-nothing threshold improves financing efficiency. While neither simultaneous nor ordered, sequential investment translate directly to reality, the models both highlight tensions that likely impact securities-based crowdfunding. How each tension impacts crowdfunding platforms and how investors actually “arrive” on crowdfunding platforms are open questions for empirical research.

While empirical research on securities-based crowdfunding is currently limited, the proliferation of crowdfunding platforms will provide a wealth of future data. Existing platforms target different types of entrepreneurs and investors, and promote different policies and practices which will allow researchers to determine which factors lead to the best financing efficiency. Lack of data on the success of funded projects will prevent judging US-based crowdfunding efficiency for some time, but platforms are already judged on a number of metrics. StartEngine and Wefunder have attracted the largest number of campaigns by focusing on entrepreneurs, and Wefunder has secured the most total financing. Campaigns have the highest success rate on NextSeed, which focuses on local retail businesses. SeedInvest raises the most per campaign by attracting accredited, high-net-worth investors and providing “highly vetted investment opportunities.”²² The different approaches, target entrepreneurs, and investor provides ample motivation for future theoretical and empirical research.

One approach that may be particularly fruitful is a focus on local ventures and local investors. NextSeed and LocalStake both focus on local companies who provide retail products or services. Brown et al. (2018) emphasizes that local companies should be paired with local investors to provide efficient crowdfunding. By targeting investors who have personal motivations for contributing, the platform can overcome the information aggregation problem highlighted by Brown and Davies (2020). The key

²² <https://venturebeat.com/2018/06/09/todays-best-crowdfunding-platforms-by-the-numbers/>.

is to target investors who derive private benefits from a venture, either through being a customer, receiving local economic benefits, or deriving pleasure from financing the project (as friends and family may). As crowdfunding campaigns cannot solely rely on friends and family, it is more natural to consider local businesses and their local investors. Anecdotally, many successfully funded securities-based campaigns have been for breweries and restaurants that cater to local clientele and provide benefits to the local economy. From a research perspective, it will be interesting to examine whether local investors and local business pair to give more efficient financing, or just more breweries per capita. If the local pairings do create better efficiency, is it due to less information asymmetry between the entrepreneurs and investors, or better aggregation of investors' information?

In addition to selecting a platform to run their campaigns, new ventures must also consider what type of security to offer. Many platforms support multiple types of securities, including preferred equity, convertible and traditional debt, revenue-share agreements, and simple agreements for future equity (SAFEs). From a theoretical perspective, Li (2018) considers the optimal financial contract in securities-based crowdfunding, showing that profit-sharing agreements can efficiently aggregate investors' private information. While information aggregation is particularly relevant in crowdfunding, ventures should also consider the ability of the security type to address the more general problem of asymmetric information between the entrepreneur and investors.

The seminal work of Myers and Majluf (1984) shows that firms will sell the least information-sensitive claims first, suggesting that revenue-share agreements may be particularly attractive. By insulating investors from business costs and other risks, revenue-share agreements may be less sensitive to asymmetric information. Furthermore, revenue is likely sensitive to consumer demand, which may be equally opaque to entrepreneurs and investors alike. Revenue may also be less subject to moral hazard and reporting concerns, as there are fewer degree of freedom in reporting revenue relative to net income. How often and when revenue-share agreements are used will be an interesting topic for both theoretical and empirical research.

Another innovative security type is the simple agreement for future equity (SAFE). SAFEs act similarly to warrants, providing investors with equity at a future date when another round of financing occurs (and is priced). SAFEs may be particularly useful when valuations are very speculative and neither entrepreneurs nor investors can price the offering, i.e., when uncertainty is high but information asymmetry is low. Future research should address how often SAFE instruments are used, how they are exercised, and what discounts

early round investors receive. More generally, it will be interesting to understand how all the potential securities are used, when they are used together, and how investors respond to each security. Moreover, how does the use of these securities vary across countries and legal settings?

2.3.4 Initial Coin Offerings

Initial coin offerings (ICOs) are an interesting blend of crowdfunding and blockchain technology. Rather than pre-selling products or selling pure securities, ICOs offer coins or tokens whose ownership is recorded on a blockchain. The coins or tokens can represent access to a service, such as a computing platform, can represent control or cash flow rights like a security, or can represent some combination of the two. Typically, the services offered involve some type of technology infrastructure or networking that is only valuable when a sufficiently large network is established. Li and Mann (2019) show that by pre-selling tokens, ICOs are able to overcome several coordination failures that may result if tokens are not pre-sold. Thus, an ICO can aggregate information regarding network products, making one aspect of an ICO similar to a product pre-sale.

Many of the current questions surrounding ICOs have to do with regulation. For example, should ICOs be regulated as security offerings? When the coins or token contain cash flow rights, control rights, or are exchangeable for cash, regulation seems much more likely. Practitioners and researchers are also interested in how to avoid fraud in ICOs. In 2017, 80% of ICOs (by number) were scams, with the largest raising over \$600 million.²³ How to detect and prevent fraud in ICOs remains an open question. Finally, if ICOs function like product pre-sales, are the success factors for ICOs similar to those for rewards-based crowdfunding? Or alternatively, do ICOs behave more like securities-based crowdfunding? If ICOs are more similar to traditional financing, do future offerings and dilution affect ICO offerings similarly? Answering these questions may help to improve future offerings and guide future regulation.

Given the wide variety of offering methods that entrepreneurs now have, it is unclear how they choose the best early-stage financing option. While the business type and product clearly dictate the choice space, a more nuanced theory may help to shine further light on the issue. Regardless of further theoretical developments, relatively little work has documented how entrepreneurs are choosing to list securities. Furthermore, when an

²³ <https://cointelegraph.com/news/new-study-says-80-percent-of-icos-conducted-in-2017-were-scams>.

entrepreneur selects a particular method, we know little about how they pick a platform.

3 Personal Finance

A fundamental area of finance research addresses how individuals should allocate their wealth among available investment opportunities. Research shows how assets can be combined to create efficient portfolios, how diversification can reduce idiosyncratic risk and how risk-sharing and hedging can improve investors' welfare. While FinTech does not alter these established results, it does provide new tools that may allow investors to create more efficient portfolios, better diversify their portfolios and hedge the individual risks they face. Moreover, FinTech may lead to better education of investors and more customized solutions for individuals, ultimately enhancing investors' financial planning abilities. In the following sections, we discuss new innovations in investments, how they can change the options available to investors, and relevant research questions that may be addressed.

3.1 New Investment Opportunities

In creating new means of providing capital to individuals and new ventures, FinTech simultaneously creates new investment opportunities. Individuals can now invest in personal debt, new start-up companies, cryptocurrencies, and hybrids of these, such as Initial Coin Offerings. Importantly, these new investment opportunities, unlike hedge funds, private equity, and venture capital, are typically open to all investors, regardless of their wealth.²⁴ As a result, all investors have access to new assets that may provide return characteristics that allow for more efficient portfolios. However, the ability for these assets to truly improve investors' portfolios depends on the degree of adoption by investors and the assets' realized return characteristics.

Investment platform fragmentation is a large hurdle to the wide-spread adoption of new investment opportunities. As of 2018, there are thousands of crowdfunding platforms, over 1,600 cryptocurrencies, and over 500 cryptocurrency trading platforms. While many investors may be interested in participating in a new crowdfunding offering or buying a cryptocurrency,

²⁴ YieldStreet provides investors access to asset classes that traditionally have only been accessible to high net worth investors, such as art, real estate, commercial shipping, etc.

the myriad options available are overwhelming. And although each investment opportunity may attract some individuals, the value in many platforms or cryptocurrencies is increasing in the number of members. This suggests consolidation may greatly help the adoption of these new investment opportunities, however that consolidation is achieved.²⁵ Large asset managers may be able to purchase platforms and market them to their existing customers. Alternatively, new business may emerge to provide curated access to new investment opportunities. By bundling assets from crowdfunding, peer-to-peer lending, and cryptocurrency platforms, new businesses may be able to give investors exposure to these new assets without the complications of dealing with each platform and investment opportunity separately. Moreover, the businesses may be able to screen companies and individuals, finding the best investments for their clients. Taken to scale, businesses may even be able to create exchange-traded funds (ETFs) that provide access to these alternative asset classes or synthesize returns with similar properties.

While the potential to intermediate new investment opportunities may alleviate search frictions and improve diversification for individual investors, the intermediation itself may make the investment opportunities less attractive. Most directly, investors would almost certainly pay an asset management fee to the intermediaries, reducing the investment return. Indirect effects may also be felt on the individual platforms. In crowdfunding, aggregating the wisdom of the crowd depends on individuals' making the investment decisions. If the "decisions" of many investors are aggregated and made by one intermediary, the power of the crowd in guiding investment decisions is mitigated. As a result, financing efficiency and investors' returns may suffer.

In addition to search and information frictions, new digital assets face costs associated with information storage and protection. Digital wallets, which can be physical hardware or software that resides locally or online, are available to store investors' digital assets and passwords. While these wallets improve the security of digital assets, hackers have strong incentives to break through security measures, which may lead to continuously changing digital security. As holding and storage costs clearly impact many commodity goods (although those costs often scale linearly with the amount of the asset, such as oil, unlike digital assets), how will storage and security costs impact the expected return on digital assets? How will those costs and the types of security evolve, and will investors have to constantly pay attention in order to stay protected?

²⁵ When the same asset is traded across multiple platforms, one means of consolidation could be based off Regulation National Market System, which coordinates trading across platforms with the intention of assuring best price executions for investors. However, given that many new platforms are in different countries, there are likely many regulatory hurdles due to differing regulations regarding crowdfunding and cryptocurrencies.

While new investment opportunities clearly come with some challenges, it remains to be seen whether the investments' return properties are worth the effort to investors. Given the novelty of these products, it is too early to say in many cases, but it does raise interesting questions. Do the assets provide unique returns that help to diversify investors' portfolios and push out the efficient frontier? How similar are the returns to those in venture capital or private equity? Even if they do have similar return properties, providing easy access could improve investors' welfare. Alternatively, democratizing access to asset classes such as hedge funds and private equity could provide another business opportunity and have a similar impact on welfare. Furthermore, how should we expect competition among investors to affect the returns? Will high demand for the assets drive down returns to an unattractive level? Documenting the return performance of new investment opportunities will be critical as they develop.

3.1.1 Cryptocurrencies

The spectacular rise and volatility of Bitcoin captures news headlines and stokes investor excitement, but that same roller-coaster movement also leads critics to doubt its viability to become a cash equivalent. The promise of a decentralized system to replace our existing monetary system (which is regulated and managed by federal banks) has piqued interest since the early days of the Internet; Milton Friedman said in 1999 that “a reliable e-cash” “will soon be developed” and will become “one of the major forces for reducing the role of government.”²⁶ It is not yet clear, however, whether cryptocurrency in general, or Bitcoin in particular, will be able to achieve that goal. Several studies document price manipulation in cryptocurrencies (Griffin and Shams 2020; Gandal et al. 2018), leading to concern that manipulation may destroy confidence in cryptocurrencies and prevent wide-spread adoption. Even if price manipulation can be mitigated, some observers lament that due to its highly technical nature, cryptocurrencies have a utopia problem.²⁷ Cryptocurrencies can work very well if the average Joe and Jane trust them to the same extent as government-based currencies, and know how to use them in the same way they know how to use cash, but such a standard is currently impossible.²⁸ Cryptocurrencies rely on their own ecosystems, which

²⁶ <https://www.coindesk.com/economist-milton-friedman-predicted-bitcoin>.

²⁷ <https://medium.com/livenpay/the-utopian-problem-c703dfab1e4>.

²⁸ To be used as cash, cryptocurrencies must also be transferred and settled quickly. However, cryptocurrencies often rely on Proof-of-Work protocols that delay block creation. Saleh (2021) shows that Proof-of-Stake protocols are an economically viable alternative.

in turn rely on substantial sets of believers. In other words, cryptocurrencies have an even higher level of difficulty for diffusion than traditional products (Rogers 2010). The potential competition among different cryptocurrencies, each with their pros and cons, is certainly another area where much research is needed.

In addition to the adoption problem, cryptocurrency will sooner or later collide with the current government-mandated currency system. Federal banks around the world still look at cryptocurrency skeptically. In the United States, former federal reserve chair Janet Yellen expressed concerns over cryptocurrencies in 2017: “it is not a stable store of value, and it doesn’t constitute legal tender.”²⁹ China banned cryptocurrencies and is even considering banning bitcoin mining altogether.³⁰ Given these realities, there remain significant challenges for the future growth of cryptocurrencies, regardless of how fast their values grow among traders. More research, and perhaps further developments for cryptocurrencies (e.g., some middle ground between traditional monies and the current Laissez-faire approach of cryptocurrencies) would be needed to make a viable “e-cash” system to balance the pros and cons of each scheme.³¹

In addition to cryptocurrencies, another interesting but less eye-catching development is local currencies. The Sardex currency, a B2B regional credit network originating from the Italian island of Sardinia, is one example. It has grown remarkably well, but it remains an open question whether it will scale up over larger geographical areas.³² More importantly, it is not exactly a decentralized version of e-cash because businesses must go through a vetting process before enrolling. What will be the role of such hybrid forms of “cash” in the future?

3.2 Robo-Advisors

Many FinTech businesses focus on established asset classes and provide individuals cheaper or customized means to construct portfolios and manage their investments. These robo-advisors vary substantially in the products and services they offer, but the common thread is that they provide investment options at a fraction of the traditional cost with less direct human

²⁹ <https://www.cnn.com/2017/12/13/fed-chief-yellen-says-bitcoin-is-a-highly-speculative-asset.html>.

³⁰ <https://www.reuters.com/article/us-china-cryptocurrency/china-wants-to-ban-bitcoin-mining-idUSKCN1RL0C4>.

³¹ See Thakor (2019) for an in-depth discussion of the issues surrounding cryptocurrencies and fiat money.

³² <https://www.ft.com/content/cf875d9a-5be6-11e5-a28b-50226830d644>.

intervention or advice. As an example, Betterment provides personalized asset allocation recommendations, access to globally-diversified portfolios, automatic rebalancing to target asset allocations, and automated tax-loss harvesting for 25 basis points per year. Other services also offer socially responsible funds (Wealthsimple), access to financial advisors (SoFi Wealth), or similar services for companies' 401(k) plans (Bloom). In each case, individuals are promised more efficient asset management at a fraction of the cost.³³

The growth of early robo-advisors has led many of the traditional asset managers to lower fees and either acquire or develop their own robo-advisor services. Vanguard, Charles Schwab, and TD Ameritrade all have their own platforms and have aggressively lowered advisory fees in recent years to avoid losing customers to the new challengers. Moreover, many fund managers have also lowered asset management fees as low-cost index funds and ETFs have gained in popularity. It will be interesting to observe how long the price war will continue, and how many firms can survive in a low-fee regime. It is also interesting to consider how the fees between advisory services and asset managers relate to one another—does lower advisory fees allow for higher asset management fees, or do shrinking fees lead to more savvy customers, leading high fee providers to lose capital?

A major factor in the rise of robo-advisors has been the increased transparency around financial advisory and asset management fees. Marketing campaigns have educated consumers regarding the costs they pay through lower returns and what services they should expect to receive. Web-based tools allow for easier price comparison between funds or services and more savvy customers. As a result, advisors, and managers may have to be even more transparent moving forward. Many funds are still sold through broker channels, in which consumers must go through a financial advisor to invest in those funds. Those funds typically have higher fees, to compensate the referring brokers. Increased transparency and the access provided by robo-advisors may enable some funds to cut out the intermediary. Will such brokered funds survive in an era of increased transparency? Will brokers' revenue model or fee structure adapt to this new environment?

While current robo-advisors primarily focus on providing efficient portfolio management, customer data is available to offer more customized portfolio management. To date, most customization is centered around investors'

³³ Robo-advisors can also help to mitigate investors' behavioral biases (D'Aunton et al. 2019).

wealth, age, and risk tolerance.³⁴ However, many other characteristics may be particularly important in forming portfolios. An investor may want to hedge their income risk by reducing investment in their employer's industry, or even by investing more heavily in competitors. Alternatively, investors may want to invest less in local firms as their income and home equity are highly correlated to local economic conditions. Will robo-advisors be able to efficiently construct portfolios for each individual investor? Or will specific ETFs be designed for employees of particular companies or industries?

A related issue is that homeowners are often under-diversified because most of their real estate exposure is through their own home. QuantmRE is offering a means for homeowners to diversify the risk in their own homes' equity. By selling a portion of the equity in their homes using smart contracts, individuals can diversify their real estate exposure and prevent their portfolios from being dominated by one asset in one specific location.

While generic solutions, such as tailored ETFs, may be an improvement, more customized solutions may better fit investors' needs. However, more customized solutions also require more information and potentially more investor education. For example, how sensitive is an individual's income to their company's stock performance and how does this change throughout their career? Are some jobs much less sensitive to the employer or industry because they are highly transferable? Better understanding the various interactions of income, local economic performance, and various asset returns may be particularly valuable as it becomes less costly to truly customize portfolios. Broadly speaking, what information about an individual is important in determining her optimal portfolio?

A clear cost of customized portfolios is the increased trading needed to establish and dynamically adjust positions. To make this cost feasible at scale for many investors, trading needs to become more efficient and less expensive. Many exchanges, brokers, and clearing operations are currently exploring blockchain-based solutions to improve trading, as well as the back-office processes of clearance and settlement. One of the most important aspects of blockchains is in coordinating trust among counterparties. By each party trusting the blockchain and its network structure (rather than each individual counterparty and business partner), counterparty risk can be reduced, which can be particularly valuable in cross-borders trading and settlement. Counterparty risk can be further reduced for contracts like futures or forwards by using smart contracts to automatically settle positions at set

³⁴ A notable exception is Ellevest, which focuses on closing a number of "gender investing gaps" that many women face, such as pay gaps, salary curve gaps, underinvestment in risky assets, and longer life spans.

dates. Blockchains also have the potential to streamline processes and reduce the time and cost of trade clearance and settlement. By reducing counter party risk and standardizing processes, trading can become easier and less expensive, which can encourage more trading and ultimately better resource allocation.

In addition to improving the back end of trading, trading itself is now being conducted entirely via blockchain. In early 2019, tZero securities began trading between accredited investors, eliminating the middlemen such as brokers and clearing agencies. Blockchain trading could give rise to customized markets for private equities, where only authorized participants could transact, ultimately making private equity more liquid. While blockchain has the potential to greatly reduce the costs of trading, if the blockchain is too slow, investors may resort to paying transaction fees to get their transactions processed into blocks sooner.³⁵ Transaction fees would add explicit costs to blockchain trading potentially offsetting the gains from other blockchain efficiencies.

3.3 Financial Planning

Financial planning goes beyond portfolio management and asset allocation and takes a holistic view of individuals' and families' financial wellness. Traditionally, individual advisors who get to know their clients on a personal level perform financial planning. Today, many businesses are offering components of the financial planning process through alternative means. For example, one of the most basic functions of financial planning is budgeting and measuring net worth. Websites such as Mint and Acorn aggregate users' various accounts to provide budgeting tools to guide spending and promote saving. In addition to new businesses that directly provide services to individuals, many financial planners are using new tools and technologies to more efficiently serve their clients.³⁶

A key aspect of financial planning is understanding a client's financial wellness from all aspects. While this is partially accomplished through direct interactions, significant effort is often required to aggregate financial data from the many accounts most individuals have. For example, many families have banking accounts, investment accounts, multiple retirement accounts (perhaps from current employers, past employers, and individual retirement

³⁵ <https://99bitcoins.com/bitcoin/fees/>.

³⁶ Pefin provides an AI system that advisors can give to clients to automate financial planning, and Wealthfront gives clients automated financial planning through Path.

accounts), real estate investments, etc. Technology enables new ways to share such information more securely and more efficiently. In the healthcare space, Timicoin aims to make patient data available via blockchain for efficient sharing across providers. The patient controls the release of the data to various “consumers” such as doctor offices or pharmacies.

Blockchain technology could also be applied to individuals’ financial information. An individual could link their financial accounts together and then control access to their data, passing it to potential service providers whom they choose. For firms, this could make the regulatory burden of “know your customer” less onerous, and for individuals, easier information transfer could result in better service in a number of ways. By having a holistic view of an individual’s finances, more precise recommendations could be made. Furthermore, decreasing the cost of information processing could greatly improve competition as firms will be more worried about losing clients to competitors.

3.4 Insurance

Insurance, which is one aspect of financial planning, helps individuals to hedge risks to themselves (health, life, home, and auto), their ability to produce income (disability), and their longevity (annuities). Technology is impacting the insurance industry in a number of ways, and within the industry, the impact of FinTech is referred to as InsurTech. As in many industries, technology is changing the ways consumers shop and receive customer service. While using phone apps to process claims is relatively straightforward, technology is also changing the format of the product. For example, several start ups offer pay-as-you-go car insurance, and you can buy insurance for as little as one hour. Data technologies are particularly valuable to insurance companies, as it opens up new opportunities in pricing and managing risks.

Increased data availability has the potential to help and harm individuals. Better information can lead to better understanding risk and better monitoring of individuals’ habits. By better understanding risk, companies may be able to offer products to individuals who were previously uninsurable due to diseases or conditions that are not actually as severe as once thought. Through monitoring, such as by recording driving habits (through mobile phones or dongles attached to the OBD-II ports of most automobiles), insurance companies can give discounts to less risky individuals and implicitly price discriminate based on behaviors. Such monitoring could potentially extend beyond driving habits, as many individuals now wear health monitoring wrist bands and have their homes connected to the Internet.

While extensive monitoring by insurance companies and increased data availability clearly raises issues of privacy, the implicit price discrimination that accompanies increased data raises interesting questions as well. Will monitoring favor higher income individuals, as they are more able to afford the monitoring technologies? Will insurers monitor behaviors or screen traits that may be typical of a protected class of people? As we get more information on people, when should people be insured? As an extreme example, should infants be insured against results of genetic screening tests that may affect future health care costs? If not, should insurers be able to use information in genetic tests to price policies? A finance solution could involve buying the option to purchase insurance at pre-determined prices, allowing individuals to insure against the risks of future information revelation.

A typical insurance problem is asymmetric information, in which an individual may be more likely to apply for insurance when they are more likely to need it. Insurers protect against such adverse selection by using extensive personal and medical information to initially price insurance contracts (when a contract is offered at all). Because processing and collecting information is so costly, these contracts are often large, long-term, and non-salable, reducing investors' abilities to adjust to changing life circumstances. A potential upside to increased information availability is that existing life-insurance policies and annuities may be more dynamically priced. As information becomes less expensive and more precise, information asymmetry becomes less severe and it may be possible for secondary markets for insurance products to develop. This could greatly improve investors' abilities to dynamically hedge risks as their situations change. By offering smaller contracts, and salable contracts, overall insurance adoption may be higher and more individuals may better hedge risks. As an example, Blueprint Income now offers customized annuities that can be sold piece-by-piece over time, essentially creating a subscription plan for annuities.

For better information to lead to a secondary market for insurance, it will be important to consider how finely individuals should be partitioned to create well-functioning secondary markets. Too broad partitions may lead to too much information asymmetry, while too narrow partitions may lead to very little liquidity in the secondary markets. As information availability improves, and it becomes less costly to verify that information (as is important in underwriting, which is often critical before life-insurance transactions), secondary markets for insurance products may become more common.

One particular threat to insurers is P2P or crowdfunding insurance, which creates risk-sharing pools for various needs. The idea is that a group of people

insures one another, rather than an insurance company serving as an intermediary. These pools are most commonly seen for medical insurance, but applications to retirement savings, disability insurance, or life insurance are also possible. Some new startups are even combining the idea of risk sharing pools with insurance. Lemonade groups customers together and pays any funds in excess of their claims to a charity of their choosing.³⁷ Current insurers may need to adopt such hybrid strategies, particularly if individuals are able to communicate and organize more easily via financial technology.

4 Concluding Remarks

This chapter reviews a broad range of recent developments in the FinTech area, with the hope of identifying opportunities for practitioners, researchers, and policymakers. FinTech is an exciting area. It has the potential to benefit society, business, and consumers, but can also be fraught with risk. Only synergetic collaborations among researchers, practitioners, and regulators can ensure that we exploit its rich potential while protecting consumers and retail investors and improving social welfare.

Acknowledgments We would like to thank Scott Cederburg, Shaun Davies, Keong Kim, Raghu Rau, Rachel Sanborn, Rick Sias, Chong Wang, Zaiyan Wei, and Jun Yang for their helpful insights and suggestions.

References

- Brown, David C., and Shaun W. Davies. 2020. Financing efficiency of securities-based crowdfunding. *The Review of Financial Studies* 33 (9): 3975–4023.
- Brown, David C., Shaun W. Davies, and Mingfeng Lin. 2018. Equity crowdfunding: Theory and evidence. Working Paper.
- Butler, Alexander W, Erik J. Mayer, and James Weston. 2020. Discrimination in the auto loan market. Working Paper.
- Chava, Sudheer, Nikhil Paradkar, and Yafei Zhang. 2017. Winners and losers of marketplace lending: Evidence from borrower credit dynamics. Working Paper.
- Chemla, Gilles, and Katrin Tinn. 2020. Learning through crowdfunding. *Management Science* 66 (5): 1783–1801.
- Cong, Lin William, and Yizhou Xiao. 2019. Information cascades and threshold implementation. Working Paper.

³⁷ <https://www.ft.com/content/db833e5a-6eb1-11e6-a0c9-1365cc54b926>.

- Cornaggia, Jess, Yifei Mao, Xuan Tian, and Brian Wolfe. 2015. Does banking competition affect innovation? *Journal of Financial Economics* 115 (1): 189–209.
- D'Aacunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi. 2019. The promises and oitfalls of robo-advising. *The Review of Financial Studies* 32 (5): 1983–2020.
- Duarte, Jefferson, Stephan Siegel, and Lance Young. 2012. Trust and credit: The role of appearance in peer-to-peer lending. *The Review of Financial Studies* 25 (8): 2455–2484.
- Ellman, Matthew, and Sjaak Hurkens. 2019. Optimal crowdfunding design. *Journal of Economic Theory* 184, 104–939.
- Freedman, Seth, and Ginger Zhe Jin. 2008. Do social networks solve information problems for peer-to-peer lending? Evidence from Prosper.com. Working Paper.
- Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery. 2019. The role of technology in mortgage lending. *The Review of Financial Studies* 32 (5): 1854–1899.
- Gandal, Neil, J.T. Hamrick, Tyler Moore, and Tali Oberman. 2018. Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics* 95: 86–96.
- Gao, Qiang, Mingfeng Lin, and Richard Sias. 2020. Words matter: The role of texts in online credit markets. Working Paper.
- Griffin, John M., and Amin Shams. 2020. Is Bitcoin really untethered? *The Journal of Finance* 75 (4): 1913–1964.
- Hildebrand, Thomas, Manju Puri, and Jörg Rocholl. 2016. Adverse incentives in Crowdfunding. *Management Science* 63 (3): 587–608.
- Iyer, Rajkamal, Asim Ijaz Khwaja, Erzo F.P. Luttmer, and Kelly Shue. 2015. Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62 (6): 1554–1577.
- Li, Jiasun. 2018. Profit sharing: A contracting solution to harness the wisdom of the crowd. Working Paper.
- Li, Jiasun, and Mann, W. 2019. Digital Tokens and Platform Building. Working Paper.
- Liberti, José María, and Mitchell A. Petersen. 2018. Information: Hard and soft. *Review of Corporate Finance Studies* 8 (1): 1–41.
- Lin, Mingfeng, and Siva Viswanathan. 2015. Home bias in online investments: An empirical study of an online crowdfunding market. *Management Science* 62 (5): 1393–1414.
- Lin, Mingfeng, Nagpurnanand R. Prabhala, and Siva Viswanathan. 2013. Judging Borrowers by the company they Keep: Friendship networks and information asymmetry in Online peer-to-peer lending. *Management Science* 59 (1): 17–35.
- Lin, Mingfeng, Richard W. Sias, and Zaiyan Wei. 2018. The survival of noise traders: Evidence from peer-to-peer lending. Working Paper.
- Mayer, Erik J. 2018. Big banks, household credit access, and economic mobility. Working Paper.
- Michels, Jeremy. 2012. Do unverifiable disclosures matter? Evidence from peer-to-peer lending. *The Accounting Review* 87 (4): 1385–1413.

- Myers, Stewart C., and Nicholas S. Majluf. 1984. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of financial economics* 13 (2): 187–221.
- Pope, Devin G., and Justin R. Sydnor. 2011. What's in a picture? Evidence of discrimination from Prosper.com. *Journal of Human Resources* 46 (1): 53–92.
- Ravina, Enrichetta. 2019. Love & loans: The effect of beauty and personal characteristics in credit markets. Working Paper.
- Rice, Tara, and Philip E. Strahan. 2010. Does credit competition affect small-firm finance? *The Journal of Finance* 65 (3): 861–889.
- Rogers, Everett M. 2010. *Diffusion of innovations*. New York: Simon and Schuster.
- Saleh, Fahad. 2021. Blockchain without waste: Proof-of-stake. *The Review of financial studies* 34 (3): 1156–1190.
- Schwienbacher, Armin. 2018. Entrepreneurial risk-taking in crowdfunding campaigns. *Small Business Economics* 51 (4): 843–859.
- Stein, Jeremy C. 2002. Information production and capital allocation: Decentralized versus hierarchical firms. *The Journal of Finance* 57 (5): 1891–1921.
- Strausz, Roland. 2017. A theory of crowdfunding: A mechanism design approach with demand uncertainty and moral hazard. *American Economic Review* 107 (6): 1430–1476.
- Tang, Huan. 2019. Peer-to-peer lenders versus banks: Substitutes or complements? *The Review of Financial Studies* 32 (5): 1900–1938.
- Thakor, Anjan V. 2019. FinTech and banking: What do we know?" *Journal of Financial Intermediation*.
- Vallee, Boris, and Yao Zeng. 2019. Marketplace lending: A new banking paradigm? *The Review of Financial Studies* 32 (5): 1939–1982.
- Wei, Zaiyan, and Mingfeng Lin. 2016. Market mechanisms in online peer-to-peer lending. *Management Science* 63 (12): 4236–4257.
- Yermack, David. 2018. FinTech in Sub-Saharan Africa: What has worked well, and what hasn't. Working Paper.
- Yermack, David. 2019. Blockchain technology's potential in the financial system. In Proceedings of the 2019 Financial Market's Conference.
- Zhang, Juanjuan, and Peng Liu. 2012. Rational herding in microloan markets. *Management Science* 58 (5): 892–912.



Innovations in Alternative Finance in Historical Perspective

David Chambers, Rasheed Saleuddin, and Craig McMahon

Technological change is rapidly disrupting the financial services sector. Start-ups and traditional banks offer a range of products and transaction convenience that were unthinkable even a few years ago. Moreover, thanks to technology-enabled cross-border payments, alternative data in credit reporting and digital savings accounts, financial services reach more customers than ever. In the near future, the World Bank expects that an

We are grateful for suggestions from Raghu Rau, Carsten Burhop, Duncan Needham, Alain Naef, Seung Woo Kim, one anonymous referee and participants at the 2017 Economic and Business History Society annual conference, 2017 Association of Business Historians annual conference, and seminars at the University of Cambridge. All errors are ours. This chapter contains text from a previously published article, see Chambers, D., Saleuddin, R., and McMahon, C. 2019. 'Alternative Finance: a historical perspective', *Financial History Review* 26(2):109–126.

D. Chambers (✉)

Judge Business School, Cambridge University, Cambridge, UK

e-mail: d.chambers@jbs.cam.ac.uk

Center for Economic and Policy Research (CEPR), Washington, DC, USA

R. Saleuddin

Cambridge University, Cambridge, UK

e-mail: r.saleuddin@jbs.cam.ac.uk

C. McMahon

Villanova School of Business, Villanova University, Villanova, PA, USA

e-mail: craig.mcmahon@villanova.edu

additional 1 billion low-income adults will have access to adequate payment services and bank accounts. (World Bank 2019). Tellingly, new financial channels such as online consumer lending, fund-raising platforms, cryptocurrencies and microfinance have developed outside of traditional banks and capital markets. Britain is at the forefront of financial and technological innovation, a fact at least partially attributable to the UK's financial regulatory 'sandbox' system which has received widespread acclaim and global support. The regulatory sandbox encourages financial innovation under a set of rules outside of existing consumer protection regulation. As of May 2019, over a hundred companies operated under its 'light-touch' purview. The fact that Britain's Financial Conduct Authority (FCA) recently admitted a fifth sandbox cohort which includes distributed ledger technology firms—a key technology underlying cryptocurrencies—indicates their interest in analysing how these markets operate. Financial firms and regulators are concerned with the same questions: how to stimulate financial innovation while offering ample consumer protection, and at what point the need arises for more formal regulatory mechanisms. These questions may be even more important for developing countries and IGOs, as it is believed that financial inclusion through technology may provide a route out of poverty, increased female empowerment and enhanced support for the SME sectors that may, in turn, lead to more robust economic growth in depressed areas and states (World Bank 2019). London is a leading centre—alongside Silicon Valley and New York—for these nascent companies, and many industry leaders and policy-makers look first to London for guidance on how to structure regulatory responses to financial innovations.

But what about earlier technological innovations in finance? Could they shed some light on how fintechs, techfins, cryptoassets and other non-traditional financial products and services might evolve and how policy-makers should respond to the resultant new opportunities and threats? Following on the heels of the Dutch Republic of the seventeenth century, Britain emerged as the leading financial centre in the eighteenth and nineteenth centuries. New alternative financial products and firms emerged and prospered alongside traditional banking and finance.

In this chapter, we analyse four periods in British and Irish history where private innovation created financial products and services that developed outside of the existing banking and capital market structure.¹ These cases

¹ This is not to say that there were not other innovations in other countries and in other time periods. For example, private currencies have likely existed since before recorded history, and were common in many British colonies, China throughout history (up until the twentieth century), the US, especially after the Civil War, Africa and in earlier times in Britain. See Lurvink (2014).

highlight instances where new industries allowed citizens of all income levels to better access markets in traded goods (by providing much-needed media for exchange), gave the poor access to consumption-smoothing loans (albeit at high cost), and granted smaller firms access to equity financing that was previously impossible. Access to SME financing, consumer credit and tools to make efficient market transactions are all considered in modern times to be conditions precedent for growth through financial inclusion (UNCDF 2018). Countries such as China have used financial innovation emerging from the tech sectors to leapfrog the traditional banking sector to allow individuals and SMEs access to new payments systems (increasing the efficiency of market transactions), peer-to-peer lending and crowdfunding (Wang and Dollar 2018). Across the globe, fintech companies are reducing transaction costs and other frictions for money transfers, banking and investing.

Our four historical cases demonstrate how entrepreneurial and innovative financial businesses responded to unmet consumer and firm demand. In each case, we document how these new markets emerged free of regulatory scrutiny and why the government subsequently responded to correct market failures: perceived, actual or anticipated. As our goal is to highlight parallels with today's nascent markets, the four cases herein were chosen based on their similarity to certain modern-day financial innovations. The first case looks at unsecured consumer lending from the late nineteenth century onwards, analogous to today's internet and app-based 'pay-day' lending platforms. Both then and now, many low-income consumers were unable to access loans through the traditional banking system. While today's entrepreneurial lenders utilize cutting-edge technology, they, like the Victorian innovators, were responding to unmet consumer demand. The second case draws insights from the junior IPO markets in the early twentieth century and invites parallels with Initial Coin Offerings of the present. Then as now, the IPO market for the raising of smaller funds began without state oversight. Yet consumer protection concerns in both cases drove new and necessary regulation.

Then as now, the market for the raising of smaller funds in innovative ways began without regulation, state or otherwise. Yet consumer protection concerns in both cases drove new and necessary regulatory responses. The third case describes the development and growth of eighteenth century privately issued token coins, arguing that they were remarkably similar to the current generation of cryptocurrencies, such as Bitcoin. Private industry, then as now, applied new technology—in the eighteenth century it was the steam press to print copper token coins – to solve what is now understood to be a significant state failure, the 'want of small change'. Today, cross-border micropayments, especially, are costly to effect, so crypto start-ups are seeking

to lessen the expense and delivery time for small value remittances between countries. Our fourth and final case deals with a long-lived experiment in what is now termed microfinance. While the lending of small sums by relying on a borrower's social capital is often considered a modern innovation, it was previously used to dramatic and widespread effect in Ireland beginning in the eighteenth century.²

While the regulatory response in all four cases was initially facilitative, there is always the possibility that entrenched interests will restrict innovation. In the fourth case, concerning Irish microlending, even though the initial state response was supportive, the commercial bank lobby later demanded and received state protection that virtually eliminated the microfinance industry. As such, we have included this shorter example as a warning to fintechs as to how the government can stifle financial innovation if a new technology threatens entrenched interests.

Obviously, there are important differences between today's innovations and those that have come before. Yet in certain ways the current crop of technology-driven financial products and services are following the historical paths of earlier innovations and regulatory response. For example, the four historical case studies were driven by the existence of a need unsatisfied by the financial system of the day. Moreover, it was the non-bank sector that met those demands with new products and services, while the government stood by and monitored these early developments. The parallels with today are real. In some respects, we have been 'here' before. What will be the next step in the lifecycle for today's financial innovations? History may provide some guide.

1 Unsecured Consumer Lending

In August 2018, Wonga, a major UK-based payday lender, collapsed into administration. Wonga was once considered a rising fintech disruptor with the potential for a billion-pound stock flotation. Following the Global Financial Crisis, Wonga grew rapidly and raised its profile through advertising including a football sponsorship with Newcastle United. Soon after, politicians and the media questioned the payday lending business model and were offended by its high-interest charges. In response, in 2015 the FCA investigated the sector and implemented price caps on all payday loans. As

² For an example of the belief that microfinance is a twentieth century invention, see CGAP FinDev Gateway, <https://www.findevgateway.org/topics/history-microfinance>, the website of an important think tank for development.

expected, under price controls the supply of funds contracted, and numerous lenders ceased operating. This episode was not the first time that regulators have targeted the high-cost consumer credit sector. This case study examines the origin of Britain's small-loan regulation and identifies three lessons applicable to the modern alternative financial industry, including how regulators: (1) struggled to define the alternative product, (2) attempted to draw boundaries between the alternative sector and traditional banking and (3) avoided licensing the sector. It demonstrates that while entrepreneurial financial companies grew in response to consumer demand, it was precisely during periods of market expansion that the state responded with regulation. While in the 1900s the government employed light-touch regulation, by 2015 the state had implemented heavy-handed price controls.

This case study focuses on high-cost, small value, short-term unsecured loans. Historically, this credit was known as a promissory note. Whereas savings banks and friendly societies provided financial services to workers with secure employment, moneylenders catered to a borrowing population that balanced weekly expenses on an unpredictable and limited income (Horne 1947: 232). Simply put, these alternative credit providers offered a service that traditional institutions avoided or could not service profitably. Moneylenders, and their modern counterparts, payday lenders, issue funds based on the borrower's ability to repay out of future wages. Before the 1990s, loans were repaid through a series of small regular payments, spread out over a period of many months. While these loans were considered as 'scraps from the banker's plate', moneylenders were responding to the unique borrowing needs of the working poor in urgent need of cash. The cost of these scraps was high, at times approaching 3000% APR. Moneylenders justified their rates owing to high administration expenses and the risk associated with unsecured lending. Still, regulatory activists argued that moneylenders price-gouged vulnerable borrowers. During an age of supposed Victorian free trade, what factors generated regulatory interest in this private market solution to a demand for cash?

Late Victorian moneylenders operated without regulatory oversight. In fact, the sector had inadvertently benefited from a sweeping act of credit price deregulation in 1854. In that year, reversing centuries of tradition, the British Parliament repealed the usury law. The rationale behind the repeal was to facilitate commercial transactions that were demanding larger and more sophisticated loan products. Few politicians were concerned with small value consumer credit, and yet the repeal 'brought into existence these swarms of moneylenders' (Farrow 1895). Thus, an unintended consequence

of the Government's decision to free traditional finance from interest rate restrictions was to incentivize alternative providers into the market.

In the 1890s, policymakers and the media focused on the sector's growth, high-interest rates and advertising. At the turn of the twentieth century, it is likely that professional moneylenders numbered between 4000 and 6000 operating in 8000 offices (*Moneylenders' federation manual*, 1913). Advertising was problematic for moneylenders. While it was vital to generate new business, it also drew attention towards the sector's high-interest rates. For concerned policymakers, the sheer volume of advertisements was proof enough that the sector was growing rapidly. The *Daily Telegraph* and *Standard* carried an estimated 18,000 moneylending advertisements annually ('Moneylenders', 1897). Alongside these advertisements, newspaper headlines drew attention to the often-dire circumstances of borrowers. In 1897, one newspaper reported that Simpson and Co. lent £5 16s to Arthur Slater, a railway fitter earning approximately £84 annually ('Ruined by Exorbitant Interest', 1897). Slater had borrowed the money to pay medical bills. After paying £9 12s to Simpson, a balance of £9 10s still remained. Reform advocates considered Slater's loan, and the many thousands like it, as predatory. Tellingly, with abundant advertising and market growth, the political will to restrict the sector was increasing. With support from Arthur Balfour MP (the future Prime Minister), the Moneylenders Select Committee 1897 convened to uncover the 'alleged evils' of high-interest loans.

The data suggest that loans carried an average 60% APR, typically ranging between 60 and 600% APR (MLSC 1897: 110; 1898: iii–v, 6, 3556). Over seven and half years, Fieldings Limited had issued 34,344 loans, 99.2% of which were on promissory notes with an average rate of 60% APR. Moneylenders attempted to educate the 1897 Select Committee as to the workings of their business. Isaac Samuel, a moneylender, compared small-loan lending to insurance underwriting, as both employed risk-based pricing (MLSC 1897: 2173). Pushing the analogy further, Fieldings Limited claimed that they did not charge interest, 'but an insurance premium on sold money'. Another analogy treated moneylenders as money retailers, Samuel explained, 'We sell £5 for £6 5s ... if a butcher sold 5s worth of meat and it cost him 4s that would be a shilling on the transaction; that is 20 per cent. If that was multiplied as you multiply moneylender's interest that would be 7000% APR (MLSC 1897: 2080). Like retailers, lenders' gross profit had to cover all expenses, the owner's salary and a return on capital. The evidence suggests that net income rates averaged 11%, a sum that moneylenders argued was not at all excessive (Commons debate, 21 June 1900, vol. 84).

Still, sceptical policymakers remained unconvinced of the necessity of such high-interest rates and felt that the industry's growth necessitated a regulatory response. Parliament debated legislative remedies including price controls, banning advertisements, disclosure requirements and licensing. While the particulars of the debate fall outside the scope of this study, insight can be provided by the political and regulatory problem of defining and legitimizing an alternative financial product.

If the impending regulation was to be effective, and for moneylenders to be certain of their legal status, Parliament needed to state what constituted a moneylending business. However easy this may seem, because the basic function of a moneylender is to provide funds to a willing borrower, it was difficult for the 1897 Select Committee to define how a moneylender differed from a 'proper' banker. Parliament focused on how moneylenders identified themselves to the public. Without an existing legal restriction, moneylenders were free to declare themselves a 'bank', 'finance' or 'trust company'. Parliament took offense that a moneylender's trade name could insinuate that they were a bank. Their goal was not only a way to protect the reputation of bankers, but also to ensure that any moneylending regulation did not unduly interfere with 'legitimate' commerce (MLSC 1897: 1532–1538). It was clear that boundaries were required to separate bank from non-bank lending.

This begs the question: how did Victorian policymakers define bankers? Bankers were understood as people who engaged in the activity of banking. Less formally, one witness identified that the source of capital distinguished bankers from moneylenders with the former taking deposits while the latter did not. Another remarked that 'no respectable bank ever acts oppressively or takes an exorbitant rate of interest' (MLSC 1897: 4494). However true (or untrue) that may be, Arthur Balfour showed little interest in segmenting the market by interest rates nor did he seek even a hint of reinstating a usury law. Working from the legal phraseology assigned to bankers, The Moneylenders Act 1900 defined moneylenders as, 'every person whose business is that of money-lending'. The 1900 Act also defined the sector by exclusion, that is, four specific industries were mentioned as not being moneylenders. These were pawnbrokers, friendly societies, any group granted special powers by Parliament to lend money and 'any person *bona fide* carrying on the business of banking or insurance'. Whether it was by exclusion or in comparison to banking, policymakers struggled to define their position vis-à-vis traditional banking; after all both institutions lent money to a willing borrower.

Still, a definition was in place, and from there the 1897 Select Committee debated two mechanisms to track and control the industry. The less burdensome proposal called for the creation of a moneylending registry which

was nothing more than a list of business names and locations. While the registry was meant to track the industry, it offered no control. The only listing requirement was payment of an annual fee, after which the registrant could conduct business legally. Surprisingly, many moneylenders sought a more robust regulatory system. They viewed government regulation as a road to legitimacy and sought stringent licensing, capital requirements, posted bonds, a requirement for audited financial statements and an industry standards board. Moneylenders believed a government-issued license requiring a screening process and monitoring of lender behaviour would add a much-needed dose of credibility to their profession.

As it turns out, Parliament agreed with this line of argument which is why they decidedly rejected the licensing measure. Simply put, the sector was considered unworthy to hold such a designation. It was recognized that a government license implied implicit approval of the applicant and by extension the sector. The 1897 Select Committee was adamantly against any licensing scheme as ‘some sort of approval or badge of respectability’ (MLSC 1898: 1780–1781; 1925, 1185–1190). Though certain MPs believed that moneylenders provided a needed financial service, fear of legitimizing the market was too great. Instead, the 1900 Act charged the Inland Revenue with maintaining a registry. Ironically, though Parliament sought to protect consumers, it avoided an enhanced licensing scheme because it feared sending mixed market signals concerning the legitimacy of the loan product.

Though the 1900 Act included a definition of the sector, in reality, Parliament had failed to draw boundaries around the alternative financial product. In doing so, the legislative measure failed to address the key question: How does a moneylender differ from a banker? Not long after its passage the Act’s shortcomings became apparent. In 1901, Horace Mansfield MP asked why the Inland Revenue had not prosecuted a single moneylender for failing to register (Commons debate, 18 June 1901, vol. 11). This reveals how misunderstood and ill-crafted the Act was. The Inland Revenue was charged with registering moneylenders but had no power to prosecute failure to do so. As expected, without a precise definition, it was unclear which businesses were required to register. In turn, the registry itself was of little value and the Inland Revenue did not devote sufficient resources towards it.

From a business perspective it may seem beneficial that the 1900 Act was ineffective in policing the sector. However, without clearly defined rules, judges were left to decide the merits of moneylending transactions appearing in court. As opposed to a uniformly enforced law, an individual judge could decide whether a transaction had harmed a borrower (Collard 1912: 172).

While for some regulatory interference was unwelcomed, the increased business risk of a case-by-case approach was far more problematic. After a period of renewed growth and heavy advertising, the issue of moneylending licensing was revisited in the 1920s. The Moneylenders Act of 1927 introduced a licensing requirement, but it also failed to draw boundaries around the sector. As the credit markets grew increasingly complex, this failure caused significant issues across the entire alternative consumer credit market.

By the late 1960s there was a sense of urgency about the legal status of finance houses, hire-purchase companies and instalment credit under the poorly designed 1927 Act. While banks were exempt, it was unclear which, if any, non-bank lenders required a moneylending license. It didn't matter that these new forms of credit had nothing to do with small cash sums issued to cover emergency and basic expenses: without a boundary all non-exempted consumer credit loans were at risk. Eventually, a regulatory patchwork solution exempted finance houses in 1967, and later The Consumer Credit Act 1974 provided a definitive solution. Unlike previous consumer credit licensing schemes which segmented providers by form, that is moneylender, pawnbroker, hire-purchase, bank, under the 1974 Act, financial institutions could issue all types of consumer credit under a single consumer credit licence. For the most part, this arrangement held until the FCA targeted payday lending in 2015.

Lessons can be drawn from this case study that are relevant to the alternative finance markets today. For decades Parliament struggled to define and draw boundaries around an alternative credit product. This created uncertainty and increased business risk. At first, policymakers feared sending mixed market signals concerning the credibility and legitimacy of these niche players and opted for a light-touch. As the market matured, so too did the regulatory system. The evidence suggests that once a company or sector reaches a critical mass, political interest in consumer protection and regulatory oversight soon follows. Recently, concerns over financial inclusion have shaped government policy towards fintech credit providers. No matter how compelling the innovation, finance companies are not immune from wider societal forces, especially when their products target low-income households. Moreover, policymakers often default to understanding new financial markets and products as they relate to the traditional banking sector. History demonstrates that financial regulation of all stripes is influenced by the incumbent banking system.

2 Early IPO Markets

An Initial Coin Offering (ICO) is the equivalent of an Initial Public Offering (IPO) in the cryptocurrency space. A startup company looking to create a new coin, app or service launches its ICO; investors buy into the offering, either with fiat currency or with pre-existing digital tokens in exchange for a new cryptocurrency token specific to the ICO. Perhaps most significantly, such issues are unregulated. The first ICO took place in 2013 and have risen so much in popularity that over US\$10 billion was raised from thousands of ICOs in each of the last two years. This hot market in ICOs echoes such similar episodic hot markets in early stage IPOs over the last hundred years or more in Britain. The latter provides us with a case study which may be of relevance to those concerned with the future of ICO markets today.

Britain in the shape of the London Stock Exchange (LSE) had the most developed stock market in the world by the early twentieth century (Rajan and Zingales 2003). While being subsequently overtaken by the US equity market in the middle of the century, it has nonetheless remained an important stock market to the present day. This is reflected in the more than 6000 firms that have gone public on the LSE since 1900—a total which is only surpassed by the US.

The regulatory approach to IPOs has changed dramatically over this long span of time, most particularly, in the case of the LSE's junior market—known today as the Alternative Investments Market (AIM). In contrast to the main market, the junior market has been the listing venue of choice for younger and smaller firms. From 1900 to 1929 the junior market was essentially an 'alternative finance' platform for early stage firms raising funds from retail investors. Until the mid-twentieth century UK listed equities were at least 80% owned by retail investors (Michie 1999). Hence, this market displays some similarities with equity crowdfunding and International Coin Offerings (ICOs) which have emerged more recently.

Similar to today, firms a century ago could go public on the main market or the junior market. In this earlier period, the LSE was self-regulated (and not statutorily regulated as it is today). However, only IPOs on the main market—known as the Official List—were subject to minimum listing requirements and were part of the formal capital markets in London. IPOs on the junior market, known as the Special Settlements market, were completely unregulated. Although firms had to apply to the LSE for a listing, permission was very rarely withheld provided a market maker could be found unlike in the case of the main market (see Burhop et al. 2014 for more detail). Indeed, junior market share prices were not collected and published in the LSE price

lists until 1915. In addition, prospectus disclosure was minimal and investors received little or no protection under the law (Chambers and Dimson 2009). For investors, it was a case of *caveat emptor*.

As a result, the junior IPO market in the early twentieth century was set up to provide entrepreneurs with fast access to external equity finance. Table 1 illustrates just how successful this alternative financial market was in the first half of the last century in this regard. Although we lack US IPO data for the first two decades of the last century, there were far more IPOs in London than in New York in the 1920s, 1930s and 1940s. Most of these UK IPOs were on the junior market. In the three decades up to 1929, around 1 in 4 of the junior market IPOs were start-ups, that is, newly established firms floated directly on the stock market. Thereafter, there is a marked decline in the number of start-up IPOs in the 1930s before they then totally disappear—except for a few new property development firms in the 1950s. Consistent with this drop off in start-ups, there is an accompanying rise in the average age at IPO of firms going public from less than 20 years to around 50 years. Start-ups did not reappear and the average age at IPO did not fall back to levels seen before 1929 until the launch of AIM at the end of the twentieth century.

This pattern in start-ups and firm age at IPO across the century is related to the ebb and flow in IPO regulation. From 1900 to 1929, regulation was non-existent on the junior market. In the following decade, the LSE imposed minimum listing requirements for all firms going public, including those on the junior market, in terms of firm age at IPO and profitability and dividend

Table 1 Number of IPOs on the LSE since 1900

	US IPOs	UK IPOs	Junior (%)	Start-ups (%)
1900s	–	486	63	29
1910s	–	445	80	38
1920s	297	662	100	12
1930s	105	397	100	5
1940s	141	269	28	0
1950s	447	348	0	3
1960s	2661	548	0	0
1970s	1640	267	0	0
1980s	4866	762	64	0
1990s	5202	641	47	7
2000s	2065	1175	80	26

Sources UK data from Burhop, Chambers and Cheffins (2014) for 1900–1913; Chambers (2010) for 1919–1939; Chambers and Dimson (2009) for 1940–1996; Paleari et al. (2014) for 1997–2009. US data from Simon (1989) for 1926–1939; Gompers and Lerner (2003) for 1940–1959 and Jay Ritter’s website for 1960–2009

records pre-IPO. This ultimately led to the shutdown of the junior market in 1947. Thereafter only profitable, mature firms went public and only on the main market. However, the UK IPO market became moribund in the 1970s and this contrasted sharply with the US where IPO activity was boosted by the success of the launch of its own junior market, NASDAQ (see Table 1). The NASDAQ example pushed the UK into reintroducing a junior market itself with less strict listing requirements in the shape of the Unlisted Securities Market (USM) in 1980 succeeded by AIM in 1995. This in turn led to younger firms gaining access to public equity again.

There is here a clear pattern of light-touch IPO regulation at the beginning and end of the last century coinciding with the existence of junior markets together with an intervening period of strict IPO regulation in the mid-twentieth century. However, while there are similarities between the beginning and end of the last century, the modern AIM has one important difference compared to the junior market of the early twentieth century. In the latter case, firms going public very rarely had a reputable underwriter handling the IPO (Burhop et al. 2014). In other words, poorly protected investors had nobody to certify the quality of the IPO. In contrast, firms going public today on AIM are required to appoint a nominated adviser ('Nomad'), from among a pre-screened group of boutique investment banks and corporate advisors, and it is the responsibility of the Nomad to ensure that there has been full disclosure to the market about the business and affairs of the listing firm (Gerakos and Maffett 2013, Appendix I).

While entrepreneurs clearly benefitted from getting their IPOs away, we must next consider what a lack of IPO regulation in the early twentieth century meant for investors in these IPOs. One way to answer this question is to examine firm survival (failure) rates post-IPO. Whilst estimating long-run returns post-IPO yields a more precise answer, there is considerable difficulty in accurately tracking what happens to firms that merge since merger terms are difficult to establish the further we go back in time. Although failure rate ignores the upside delivered by IPO winners, it serves to focus attention on the downside risks of IPO investing. Failure is defined as the proportion of IPOs in a given year which delist due to the disappearance or (involuntary) liquidation of the firm where shareholders receive no consideration within the five years following an IPO. Note that IPOs which subsequently disappear because of a merger with another firm and where the shareholders receive value for their shares are not regarded as having failed (Table 2).

Table 2 summarizes IPO failure rates estimated in previous studies over periods at the beginning and end of the last century and into the twenty-first century. Between 1900 and 1929 failure rates of junior market IPOs

Table 2 IPO failure rates

	Market	No IPOs	Failure rate (%)
1900–1913	Junior	561	19
	Main	264	3
1919–1929	Junior	662	20
1930–1939	Junior	397	3
1995–2010	Junior	848	10
	Main	237	3

Sources Burhop et al. (2014) for 1900–1913; Chambers (2010) for 1919–1939; Vismara and Ritter (2012), Table 5, for 1995–2010

were 20%. In other words, 1 in 5 firms going public had delisted by the fifth anniversary of going public and investors lost everything. This experience contrasts markedly with IPOs on the main market at that time where failure rates were only 3%, very similar to today, and contrasts with more recent experience on AIM where failure rates are considerably lower at around 10%. Early IPO markets also went through hot and cold periods just as they do today. There were hot IPO markets in 1911–1912 and in 1928–1929 and the failure rates of these IPO cohorts rose sharply above the average. For example, in the case of 1928–1929 IPOs, 2 in 5 subsequently failed within five years. Entrepreneurs floated some weird and wonderful schemes, and those investors who believed them suffered substantial losses. This was venture capital-type investing but in the public markets with no evidence of big IPO winners to offset the numerous failures (Burhop et al. 2014).

What becomes clear from the foregoing is that there is a relationship between IPO failure rates and IPO regulation. In the absence of regulation (and reputable underwriters), failure rates were high in the first three decades of the last century. After the poor performance of firms floated in 1928–1929, the LSE decided to steadily raise listing requirements under pressure from the government and the Bank of England, starting in the 1930s. The process of regulatory tightening continued over the following two decades. This had a dramatic effect on young firms and start-ups which were squeezed out of the public equity markets in the third quarter of the twentieth century. Since venture capital did not establish itself until the 1980s in the UK, there were no private markets able to fill this funding gap. The IPO regulatory pendulum had swung too far in the other direction. Eventually, the lighter touch regulatory approach accompanying the re-emergence of a junior market in the 1980s and 1990s has once again catered to the needs of young entrepreneurial firms wishing to access the public equity market as quickly as possible while attempting to offer minority investors protection from fraudulent schemes.

Three things become apparent from this case study of early IPO markets in the UK that are relevant to alternative finance markets today. First, regulation (including self-regulation) needs to strike a balance between providing entrepreneurs with flexible and fast access to new pools of capital, on the one hand, and protecting outside investors, on the other. Second, regulators reacting to the emergence of new markets will, in all likelihood, not get this balance right. They may under-react, or, as in this case, over-react once problems (investor losses) arise. Along the way, some institutional learning takes place. It took the LSE the best part of a century to learn from their experiences and arrive at a more appropriate balance. Third, the experiences of early IPO investors would not have mattered if these investors were ‘sophisticated’ in the sense that they understand the investment risks involved and were sufficiently wealthy to easily withstand the IPO losses. The total absence of any regulation regarding investor suitability in the early twentieth century almost certainly meant that unsophisticated investors were suckered into these IPOs. The parallels with the market for ICOs in 2018–2019 are striking.

3 Private Currency Markets in the Eighteenth Century

Satoshi Nakamoto³ founded Bitcoin to reduce transaction costs and times. More appealing to early adopters was its anonymity and inflation-proof nature (Nakamoto 2009). Advocates believe distributed ledger technology mitigates government tendencies to inflate fiat currency and offers a safe haven from the banking systems increasingly obtrusive data-collecting practices. Yet this was far from the first time that new technology was used by the private sector to cure a perceived state failure. In fact, Bitcoin and similar cryptocurrencies are simply the most recent and most technologically advanced of a wide variety of private solutions previously utilized to replace sovereign money in certain transactions. Private currencies should therefore be considered as important precursors to today’s distributed ledger currencies such as Bitcoin. Furthermore, they evolved in ways easily discernible by today’s modern fintech market participants. Successful private currencies can be classified by (1) the level of trust in the key private institutions and (2) the degree of state failure in the currency markets. The latter is evidenced by either a lack of trust in the central government or an otherwise unfulfilled need. Cryptocurrencies exhibit many of the same characteristics and

³ Likely a pseudonym.

address similar needs as the private solutions of the past. As such, analysing the history and development of alternative (private) currencies may provide useful guidance for policy related to new non-state cryptocurrencies.

Throughout history, market transactions have often relied on tokens and metallic coins as a medium for exchange of goods and services. While the sovereign state often monopolized the issuance of currency, token currencies provided by private companies or citizens were used alongside or in place of sovereign money for small market transactions at various points in history, even if such private tokens were not considered legal tender and were generally not able to satisfy tax obligations.⁴ As with sovereign-issued money, the usage of private tokens depends on trust in the issuers, or on the trust that the token will be accepted in subsequent transactions. As an alternative to currency, state-backed or private, the records of the credit resulting from marketplace transactions can be kept in (private) ledgers by merchants, temples or churches.⁵ Yet such credit arrangements have flaws that physical coins can overcome, with the latter allowing anonymity in transactions, and the flexibility to deal with those where a lack of social bonds preclude a trusting relationship.

Private token currencies developed outside the formal monetary system, and yet often dominated sovereign money for long periods. Such tokens could even circulate nationally, as they did in England over a very long time-frame, in China over many centuries but especially the long sixteenth century and in the US from time to time (Falkner 1901; Klein 1974; von Glahn 1996; Kuroda 2005, 2006; Goetzmann 2017: 387). As such, private money can surprise in geographic scope and scale, at times completely dominating the national economy (Martin 2013: 238).

Private currencies developed to supplement a lack of official coin. Indeed, Jean-Baptiste Say specifically stated that ‘there can never be any shortage of money, since the enterprising mercantile class will always be able to improvise an alternative’ (Martin 2013: 206). Parallel currencies did create some problems for users, and in this regard, private currencies were far from ideal. For example, default of the issuer was always a possibility, and in some cases a reality. As such, for a private currency to succeed, trust in the issuer is vital. Or, in cases of total state failure, private citizens need to be desperate enough

⁴ In this particular instance, we are speaking of currency mostly as a medium of exchange and unit of account. It would be rare, even if possible, for merchants and their customers to accumulate large amounts of small token coins as a vehicle for savings, due to the sheer unwieldiness of token metals such as copper compared to other alternatives.

⁵ Credit was also commonplace within smaller communities where trust was available. Yet ledgers that kept track of debts had weaknesses, such as the tendency for debtors to default. For more on private credit and private ledgers, see Muldrew (2001).

to accept any currency in transactions. Keeping track of a plethora of small change alternatives might also involve risks in exchange. However, merchants and their clientele were generally sophisticated enough to readily manage many different sovereign, private and even foreign metallic and token coins, and adjusted their prices accordingly (Kuroda 2005).

Britain—like many other countries—suffered from a shortage of small currency following the withdrawal of the Romans, especially after the Great Debasement of 1560 and the Elizabethan Great Recoinage 1696–1699, a problem known as ‘the want of small change’ (Williamson 1889). The main reason for the failure is disputed, but it does appear that the capture by various elite groups—first by landowners and creditors and then international merchants—explains a good deal of the reluctance to provide for small market transactions (Desan 2014). Whatever the reason(s) for the state failures in Medieval and early modern Britain, British enterprise has a long history of providing private currency solutions to mitigate the government failure to provide enough small denomination currency. While credit could be used for those who shared local bonds of trust, other market transactions needed a more anonymized solution. Without private enterprise, the British economy would have operated with significantly less currency, and, therefore, with less trade.

The main focus of this case is the period known as the free token era which began in the 1740s. Numismatists and historians have written extensively on the tokens of this period, including a full-length treatise by economist George Selgin (2008).⁶ The Great Recoinage of 1696 to 1699 caused perhaps the greatest ever shortage of small change, just at the time when an increasing number of new companies needed small change to pay workers migrating to the factories of Industrial Revolution Britain. While the state’s policy of strong money benefitted many elites, workers suffered greatly. Employers scoured the countryside for small coins, wasting valuable time resources, and often paying a premium for them. Most solutions were extremely unfavourable for the workers, including paying workers in groups in larger denomination coin, delaying workers’ pay, sometimes for many weeks, payment in kind, the use of stores run by the companies themselves and agreements with services such as pubs to offer credit to the otherwise unpaid workers (Selgin 2008: 24–26).

The earlier private tokens—of the 1600s—were mostly copper or brass but could also be leather or paper (Falkner 1901). They were mostly round but there was very little standardization. The innovation in the eighteenth century

⁶ Scholars include Ruding (1840), Peck (1960), Whiting (1971), and Larkin and Hughes (1973).

was, perhaps unsurprisingly, the steam engine. While earlier tokens were of differing shape, quality and even material (Falkner 1901), by 1787, accurately round and consistently weighted privately-issued 'trade' tokens, often referred to as 'Conder' tokens, became the most prevalent form of small change due to their being struck using a steam-driven press, novel use of technology that allowed for standardization and mass production at the lowest cost (Peck 1960; Selgin 2003). This new technology also made the coins difficult to counterfeit; this was true technological innovation.

While the state was unable or unwilling to use the new minting technology to grease the wheels of commerce, the private sector was, and adopted the new minting technology as soon as it became viable. Privately minted coin effectively solved the big problem of small change during this important time period, which spans both the agricultural and industrial revolutions. 'Entrepreneurs' kept the economic model in 'equilibrium, allowing the monetary system to function properly' (Sargent and Velde 2002: 302). As late as 1812, tokens were circulating at five–ten times the rate of official currency (Falkner 1901). While the private sector had provided an effective solution in the late 1780s, it took the state until 1817 to catch up with the private sector, when the copper token system was nationalized. The state, watching the circulation closely by now, nationalized the copper token system by decree in 1817 and attempted to prevent the further issuance of the private tokens. The transfer of copper token coin supply from private issuers to the sovereign mint was practically complete. State interference was fortuitous to some extent, as the private solution was not entirely satisfactory. By 1797, even the private mints could not keep up with demand. Additionally, there were frequent bankruptcies, with some issuers reneging on thousands of coins. Private failure had replaced the public failure.

The private sector had proven the worth of the copper coin as an instrumental factor in the provision of commerce before and during the industrial revolution. Indeed, the modern fiat token coin framework—as in most if not all modern currency systems—might not exist in its current form without the private solution perfected in the late 1700s (Sargent and Velde 2002: 303). Private copper mintings on a scale made possible by the steam press gave way to the copper penny we know today in the UK and elsewhere.

When taken together with other cases of the adoption of private currency or currency systems, a set of common themes present themselves. New markets that tend to develop with little intervention from the state eventually attract the state's attention. *Laissez faire* gives way to intervention. Such government regulation or control is not inevitably beneficial to all

members of society, but it equally doesn't have to have negative consequences. While banning private currencies without remedying state failures can severely impact citizens, certain government responses can provide a more robust currency compared with the purely private solution.

The state can have a tremendous impact on private currency and currency systems. It sets the stage for private competitors by failing to provide its citizens with enough usable currency, or by eroding the value of the sovereign money that is available. It can then choose how to address private competition to its sovereign rights. The state can react to private innovation in currency through an outright ban or a *laissez-faire* approach, which latter typically ends with a remonopolization. In Qing dynasty China and the US in 1674 and 1741 (Timberlake 1981), private solutions were banned, but without replacement.

Successful private currencies, on the other hand, tend to be implicitly accepted by the state for some time, as they are recognized as filling important needs that are otherwise ignored by sovereign entities. To add the Chinese example to the British one above, the Ming state accepted that temporary needs could be met by the private sector (Kuroda 2000: 191). Banning usually happens almost immediately, if it is going to happen at all. Successes generally operate in parallel with state-backed currencies. However, there comes a time when the state simply remonopolizes the supply of currency, even if it comes at great cost to certain publics, such as during the Great Recoinage. State interference has generally not been driven by meeting the needs of the public, but rather to regain control of the money supply, capture seigniorage profits, or often to reward influential elites at the expense of the majority (Desan 2014). A pattern of remonopolization is recognizable in the monetary histories of different countries. England, China and the US, often over long periods, are three important and representative examples. For example, the incoming Qing government solved the shortage by issuing more copper coins, and trade of grain between regions was outlawed. Such a government response to a shortage came at a great cost to both the state (via negative seigniorage) and certain provinces (who were unable to market surplus grain, or purchase grain that was desperately needed). Notwithstanding this great cost, the monopoly of the state on currency given up by the Ming dynasty was restored (Kuroda 2005).

Private markets have consistently identified needs that can be profitably met through innovation. Private demand drives such innovations. In this case, the private sector was the first to apply a new industrial technology—the steam press—to 'the big problem of small change' (Cipolla 1967). Once the market grew too big for the government to ignore, it simply monopolized the

technology. The government monopoly was likely to be welfare-enhancing, given that it reduced the reliance of the public on the creditworthiness of private issuers—bankruptcy risk was eliminated. As such, state responses do not necessarily destroy markets and may increase consumer welfare, as this English case tends to indicate. In the best case, private innovation often informs and then incites new state-backed currency solutions, often eliminating the private element.

The history of private currencies suggests some valuable lessons for practitioners and policymakers involved in the cryptocurrency revolution, as Bitcoin and its imitators share many characteristics with private token currencies and the centralized ledgers of the past. Firstly, it is probable that the public sector is not the best base for innovation. The private sector invented the steam press, and then used this technology to improve the lives of the citizenry who were suffering from a want a small change. Remonopolization was always a risk, yet government control did eliminate some of the frictions of the private system. Over-regulation or even banning of cryptoassets remains a strong possibility, as occurred in Britain before the free token era, and in Qing Dynasty China. However, the biggest threat to Bitcoin is likely to be a better Bitcoin, as it is far from clear that the benefits of using Bitcoin outweigh the costs (in terms of energy usage, to name one) and risks. The most likely source of competition is sovereign currencies that are reinvented to resemble cryptocurrencies to a greater or lesser extent. Once the central banks fully understand blockchain technology and its applications, they should be able to continue to dominate, or remonopolize, electronic payments. However, there are private solutions to the current limitations experienced in international and online transactions that could succeed—for example, the need for extremely small denominations (micropayments) as well as for a more efficient means for retail-sized cross-border payments.

4 Microfinance in Ireland

Outside of Nairobi, a young wholesaler has just purchased goods from a supplier. Yet instead of settling in cash or by bank transfer, the entrepreneur has applied for a loan from Musoni Microfinance. Notably, Musoni does not have any branch or retail locations. It is one of the first entirely virtual small-loan providers in Africa. Their entire loan application, processing and settlement are completed via the wholesaler's mobile phone: a truly revolutionary business model in a part of the world sorely in need of efficient capital allocation. Across Africa and many other developing countries microfinance

has experienced rapid penetration of traditional banking markets. The importance of these developments lies in the critical contribution made by access to finance for small entrepreneurs towards accelerated economic growth rates (Armendariz 2005).⁷ While Musoni and its digital competitors are a relatively new market niche, microfinance is not a modern invention. Like the other recent innovations covered by this chapter, small-scale financing to entrepreneurs had its origins long ago. For example, early microfinance was especially important in Germany, Italy and the British Isles beginning in the nineteenth century (Galassi 2001; Guinnane 1997). This case study examines how Irish microfinance institutions addressed a market failure and provided credit to the poor from the 1720s until the mid twentieth century (Hollis and Sweetman 1998a, b, 2001).

The microfinance innovation of the eighteenth century was crucial to Ireland's navigation of the Industrial Revolution, given there were no other sources of finance for small businesses (Hollis and Sweetman 2001: 292). Long before the launch of microfinance ventures in Bangladesh and parts of Latin America in the second half of the twentieth century, the author of *Gulliver's Travels*, Jonathan Swift, was providing funds to lend, interest-free, to 'poor industrious tradesmen' in the early eighteenth century (Sheridan 1787: 234). At the time, and even as late as the early twentieth century, limitations in the Irish legal system, alongside massive emigration, meant that keeping track of borrowers was extremely challenging for lenders (Guinnane 1994). Swift's innovation was to understand that the social capital existing within local communities could be utilized in order to mitigate the risks of lending to unknown entrepreneurs. This same logic underlies the activity of today's microfinance institutions. In eighteenth-century Ireland, each borrower required guarantees from two neighbours. The interest-free loans, for up to GBP 10, required weekly servicing. Though subsidizing the interest-free nature of the loans, Swift believed in 'punctual repayments', requiring financial discipline from the borrowers and their guarantors (Sheridan 1787: 234).

The Dublin Musical Society was so successful in imitating Swift that government legislation permitted the Society to expand their lending outside of Dublin in 1778. The Society used profits from its performances to subsidize loans of up to GBP 4. In 1822, the Irish Reproductive Loan Fund Institution (RLFI) provided GBP 55,000 to establish approximately 100 microlenders, again with a maximum loan size of GBP 10. Soon after, thanks in part to the RLFI's lobbying efforts, the sector secured a stamp tax break in

⁷ Such lending has not been without controversy and has attracted significant criticism in academic and development policy circles. See, for example, Bateman (2010).

1823. This change, plus permission to charge interest, incentivized competition. New microfunds, such as the Central Loan Fund Board, quickly became the primary source of credit for small businesses. Like the other case studies examined in this chapter, their success also attracted new legislation. Parliament established a regulator for all non-RLFI lenders and set a very high deposit rate cap. The latter fuelled further growth by allowing depositors to earn more on their money with the microfinance organizations than with a commercial bank. Hollis and Sweetman (2001: 1879) estimate that by 1843 there were 300 loan funds making about 500,000 loans per year, which was practically all of the credit available to the poor (Hollis and Sweetman 1998a: 353). At one point, microlenders serviced 20% of the population, none of whom were commercial banking clients (O'Grada 1994; Hollis and Sweetman 1998a, b). That is, the two industries targeted completely different clienteles, and without the Loan Funds, the poor would have no access to debt finance. The Loan Funds could be said to have found a market niche that otherwise would not have been addressed.

State support was one of the major reasons for the success of these early microfinance lenders. But aggressively prohibitive regulation was to follow. Like many private solutions, the Loan Funds were far from perfect. A government inquiry in 1896 found that some loans were predatory in nature, and fraud and overly robust recovery tactics were commonplace. Yet for many years, commercial banks could do no better in servicing the working poor, dealing as they did only with the richest third of Ireland (O'Grada 1994). The industry survived the Great Famine, attracting more avaricious lenders. However, the Board, who had since become the chief regulator, was unwilling to clamp down on the worst of the predatory behaviour. Worse still, the powerful banks began to view the loan funds as competition, and in 1843 the former were able to successfully lobby Parliament for a lowering of the deposit rate cap to a level that made the industry substantially less attractive. Elite capture is the most likely explanation for this sudden reversal in state attitude to the Loan Funds. Similarly, legislation in 1900 and 1914 put the loan funds at an even greater disadvantage and the industry essentially vanished over the first half of the twentieth century. The bank lobby had successfully wiped out a competitor.

As in the other cases featured here, private initiative resulted in a new innovation that has since been copied across the globe. Yet it was only once the microfinance initiatives had experienced significant success that they attracted the interest of regulators who were acting in the interests of the local elite. In this case, the regulation resulted in the destruction of an industry that had the potential to support smaller entrepreneurs in emerging Ireland.

5 Conclusion

In the last few years, online consumer lending, ICOs, cryptocurrencies and microfinance have all emerged and flourished. In each case, the underlying innovative idea has appeared in an earlier historical period. As this chapter argues, we can use our understanding of the evolution of these earlier innovations to provide several insights for industry participants and policymakers today. The first insight is that the private sector can often be counted on to anticipate otherwise unmet needs. Moneylenders filled a gap for consumption-smoothing loans unaddressed by the banking system of the day. Loan Funds provided debt financing for small businesses that banks in Ireland were unwilling and unable to provide. Private currencies filled a void created by the lack of sovereign coins for smaller transactions at various points in history, most notably—but not uniquely—in eighteenth century Britain. Junior market IPOs provided access to funds for smaller companies that might otherwise have struggled to raise external finance.

The second insight is that these new financial markets and instruments eventually attracts the attention of regulators. The initial regulatory response is likely to be inadequate and subsequently evolves over time. Governments may intentionally choose not to interfere, or may be ignorant of the situation, until the market grows to a point where it attracts the attention of enough influential interests who pressure the state to intervene, generally citing consumer protection grounds. The impetus for regulation may come from injured parties, the incumbent banking system, or directly from the government itself. Regardless of the motivation, it is clear that an initially *laissez-faire* or experimental regulatory system should not be taken for granted and stricter, sometimes highly restrictive, regulation is likely to follow. Financial professionals and policymakers should be aware that an initial ‘sandboxing’ and benign response to new innovations does not preclude regulators from acting when the markets hit some ‘tipping point’ in scale and scope. The case studies in this chapter are a warning to the industry to watch for shifts in the existing regulatory approach. In all four cases, harsher regulation—in some cases even to the point of shutting down the products and markets concerned—arrived after prolonged periods of observation. Tougher regulatory responses may be facilitative, even if they are harmful to select parties. State-issued copper pennies were an improvement on those issued by the Parys Mine Company, and early Irish microfinance legislation spurred growth in the sector. Yet responses from the state can be far less benign. In the cases of moneylending and microlending, the behaviour and lobbying efforts of the traditional banks influenced, and even captured, the regulatory process.

Entrepreneurs should take heed of how the existing banking community responds to financial innovation.

Finally, it is clear from history that regulation needs to strike the appropriate balance between facilitating innovation and protecting consumers if technological financial innovation is to evolve. Without an initial *laissez-faire* period, innovations allowing access to funds for business and consumers alike, such as private token coins, moneylending, microfinance and the early stage IPOs may never have evolved. On the other hand, financial products and services have been shown to be harmful at times and an appropriate level of consumer protection benefits society (Turner 2009). Regulation benefitted gullible IPO investors and protected borrowers from abusive lending practices. State remonopolization of copper token currency in the nineteenth century reduced the exposure of the users of small change to the default risk of private issuers, or to counterfeits.

Prohibition of new innovations may have negative consequences. The regulation that eventually put an end to microfinance may have resulted in a dearth of exactly the kind of financing supportive of economic growth. In the private currency case, ‘small change’ facilitated payments to the workers who were vital to the success of the industrial revolution.⁸ Today, for example, it has become clear that not providing citizens with the tools to effectively make market transactions is welfare-destroying, as it probably was in eighteenth-century Britain. The market needs to be given room by the state to innovate in a way which balances innovation with adequate consumer protection. For centuries, traditional and alternative financial providers have channelled funds to meet the demands of business and consumers. Although twenty-first-century fintechs are driving down costs and enhancing the customer experience by challenging the traditional banking sector, history tell us that there is a likelihood of the introduction of consumer protection measures that will impact these new markets and their competitive landscapes. Based on such precedent, today’s financial innovators may benefit from engaging policymakers sooner rather than later.

⁸ In the US, where there were fewer options, the lack of small change resulted in the ‘re-enslavement’ of newly freed workers in the South and may have contributed to the relatively slow economic growth in this region towards the end of the nineteenth century (Lurvink 2014).

References

- Armendariz, B. 2005. *The economics of microfinance*. Cambridge: The MIT Press.
- Bateman, M. 2010. *Why doesn't microfinance work? The destructive rise of local neoliberalism*. London: Zed Books.
- Burhop, C., D. Chambers, and B. Cheffins. 2014. Regulating IPOs: Evidence from going public in London, 1900–1913. *Explorations in Economic History* 51: 60–76.
- Chambers, D. 2010. Going public in interwar Britain. *Financial History Review* 17 (1): 51–71.
- Chambers, D. and E. Dimson. 2009. IPO underpricing over the very long-run. *Journal of Finance* 64 (3): 1407–1443.
- Cheshire Observer. 1897. Moneylenders, 15 May.
- Cipolla, C. 1967. *Money, prices, and civilization in the Mediterranean world, fifth to seventeenth centuries*. New York: Gordian Press.
- Collard, C. 1912. *The Money-Lenders Acts, 1900–1911*. London: Butterworth.
- Desan, C. 2014. *Making money: Coin, currency and capitalism*. Oxford: Oxford University Press.
- De la Torre, A., E. Levy Yayati, and S.L. Schmuckler. 2003. Living and dying with hard pegs: The rise and fall of Argentina's currency board. *Economica* 5 (2): 43–107.
- Farrow, T. 1895. *The Money-Lender unmasked*. London: Roxburghe Press.
- Falkner, R.P. 1901. The private issue of token coins. *Political Science Quarterly* 16 (2): 303–327.
- Galassi, F.L. 2001. Measuring social capital: Culture as an explanation of Italy's economic dualism. *European Review of Economic History* 5 (1): 29–59.
- Gerakos, L., and M. Maffett. 2013. Post-listing performance and private sector regulation: The case of AIM. *Journal of Accounting and Economics* 56 (2): 189–215.
- Gilbert, E., and E. Helleiner. 1999. *Nation states and money: The past, present and future of national currencies*. Routledge.
- Goetzmann, W.N. 2017. *Money changes everything: How finance made civilization possible*. Princeton University Press.
- Gompers, P., and J. Lerner. 2003. The really long-run performance of IPOs: The pre-NASDAQ evidence. *Journal of Finance* 58 (4): 1355–1392.
- Guinnane, T. 1994. A failed institutional transplant: Raffeisen credit cooperatives in Ireland, 1894–1914. *Explorations in Modern History* 31: 38–61.
- Guinnane, T. 1997. Regional organizations in the German cooperative banking system in the late 19th century. *Research in Economics* 51: 251–274.
- Hollis, A., and A. Sweetman. 1998a. Microcredit in prefamine Ireland: Explorations. *Economic History* 35 (4): 347–380.
- Hollis, A., and A. Sweetman. 1998b. Microcredit: What can we learn from the past? *World Development* 26 (10): 1875–1891.

- Hollis, A., and A. Sweetman. 2001. The life-cycle of a microfinance institution: The Irish loan funds. *Journal of Economic Behavior & Organization* 46 (3): 291–311.
- Horne, O. 1947. *A history of savings banks*. London: Oxford University Press.
- Hughes, P.L., and J.F. Larkin. 1973. *Stuart royal proclamations*. Clarendon Press.
- Klein, B. 1974. The competitive supply of money. *Journal of Money, Credit and Banking* 6 (4): 423–453.
- Kuroda, A. 2000. Another monetary economy: The case of traditional China. In *Asia Pacific dynamism 1550–2000*, ed. A. Latham and H. Kawakatsu. Routledge.
- Kuroda, A. 2005. Copper coins chosen and silver differentiated: Another aspect of the ‘silver century’ in East Asia. *Acta Asiatica* 88: 65–86.
- Kuroda, A. 2006. *Too commercialized to synchronize currencies*. Helsinki: XIV International Economic History Congress.
- Leicester Chronicle and Mercury. 1897. Ruined by exorbitant interest, 20 February.
- Lloyd’s Weekly. 1900. Interest at 400 per cent, 1 July.
- Lurvink, K. 2014. Strapped for cash: Non-cash payments on Louisiana cotton plantations, 1865–1908 1. *Tijdschrift Voor Sociale En Economische Geschiedenis* 11 (3): 123–151.
- Martin, F. 2013. *Money: The unauthorized biography—From coinage to cryptocurrencies*. Vintage Books.
- Michie, R. 1999. *The London stock exchange: A history*. Oxford University Press.
- Muldrew, C. 2001. ‘Hard food for Midas’: Cash and its social value in early modern England. *Past & Present* 170: 78–120.
- Murphy, Anne. 2009. *The origins of English financial markets: Investment and speculation before the South Sea Bubble*. Cambridge: Cambridge University Press.
- Nakamoto, S. 2009. Bitcoin: A peer-to-peer electronic cash system. Available from: <http://bitcoin.org/bitcoin.pdf>.
- Neal, Larry. 1990. *The rise of financial capitalism: International capital markets in the age of reason*. Cambridge: Cambridge University Press.
- O’Grada, C. 1994. *Ireland: A new economic history, 1780–1939*. Oxford: Clarendon.
- Orchard, D., and G. May. 1933. *Moneylending in Great Britain*. New York: Russell Sage Foundation.
- Paleari, S., D. Piazzalunga, A. Signori, F. Trabucchi, and S. Vismara. 2014. *Academic EurIPO FactBook 2014*. North Charleston: CreateSpace Independent Publishing.
- Peck, C.W. 1960. English copper: Tin and bronze coins in the British Museum, 1558–1958. Trustees of the British Museum.
- Rajan, R.G., and L. Zingales. 2003. The great reversals: The politics of financial development in the twentieth century. *Journal of Financial Economics* 69 (1): 5–50.
- Ruding, R., 1840. *Annals of the coinage of Great Britain and its dependencies* (Vol. 1). John Hearne.
- Sargent, T., and F. Velde. 2002. *The big problem with small change*. Princeton: Princeton University Press.
- Selgin, G. 2003. Steam, hot air, and small change: Matthew Boulton and the reform of Britain’s coinage. *Economic History Review* 56: 478–509.

- Selgin, G. 2008. *Good money: Birmingham button makers, the Royal Mint, and the beginnings of modern coinage, 1775–1821*. Oakland: The Independent Institute.
- Simon, C.J. 1989. The effect of the 1933 securities act on investor information and the performance of new issues. *American Economic Review* 79 (3): 295–318.
- Sheridan, T. 1787. *The life of the Rev. Dr. Jonathan Swift*. Dublin: A. Thom.
- Spufford, P. 1962. *Money and its use in Medieval Europe*. Cambridge University Press.
- Timberlake, R.H. 1981. The significance of unaccounted currencies. *The Journal of Economic History* 41 (4): 853–866.
- Turner, A. 2009. How to tame global finance. *Prospect Magazine*.
- UNCDF 2018. Igniting SDG progress through digital financial inclusion. Available at: <https://www.uncdf.org/article/3951/igniting-sdg-progress-through-digital-financial-inclusion>.
- Vismara, P., and J. Ritter. 2012. Europe's second markets for small companies. *European Financial Management* 18 (3): 352–388.
- Von Glahn, R. 1996. *Fountain of fortune: Money and monetary policy in China, 1000–1700*. University of California Press.
- Wang, W., and D. Dollar. 2018. *What's happening with China's fintech industry?* *Brookings*. Available at: <https://www.brookings.edu/blog/order-from-chaos/2018/02/08/whats-happening-with-chinas-fintech-industry/>.
- Whiting, J.R.S. 1971. *Trade tokens, a social and economic history*. Newton Abbot.
- Williamson, G.C. 1889. Historical evidence and information gathered from the traders' tokens of the Seventeenth Century and from the minor currency. *Transactions of the Royal Historical Society* 4: 171–196.
- World Bank, 2019. Universal Financial Access 2020.
- Xinwei, P. and Kaplan, E.H. 1994. *A monetary history of China*. Western Washington University.

Official Publications

- Report from the Select Committee on Moneylending 1897 (364).
- Report from the Select Committee on Moneylenders 1898 (260).
- The Moneylenders Act 1900: 63 & 64 Vic. c. 51.
- The Moneylenders Act 1911: 1 & 2 Geo. 5 c. 51.
- Report from the Joint Select Committee of the House of Lords and the House of Commons on the Moneylenders Bill 1925 (153).
- The Moneylenders Act 1927: 17 & 18 Geo. 5 c. 21.
- Consumer Credit, Report of the Committee 1971 (4596).
- The Consumer Credit Act 1974, c. 39.
- 'The Moneylenders' Federation manual and directory for the year 1913', General collection PP2501, The British Library.

Digital Currencies



A Layman's Guide to Bitcoin and Blockchain

Bhagwan Chowdhry and Seoyoung Kim

The recent, meteoric rise in crypto-activity has sparked widespread interest in this nascent asset class, which now comprises over 5000 distinct exchange-traded cryptocurrencies.¹ Worldwide google queries for the term “cryptocurrency” increased by more than fivefold in the final two months of 2017, during which time the search queries for “bitcoin” exceeded those for “trump” (see Fig. 1). There are now over 6000 Bitcoin ATMs across the world (Zmudzinski 2019),² and the average daily exchanged-traded dollar volume for Bitcoin (BTC) has tripled in the last year,³ with the latest daily activity averaging at \$27.5 billion USD for January 2020. In comparison, the average daily dollar volume for MSFT stock was approximately \$4.4 billion USD during this same timeframe.

¹ As per CoinMarketCap <https://coinmarketcap.com> on February 5, 2020.

² See <https://cointelegraph.com/news/bitcoin-atms-worldwide-hit-new-milestone-surpassing-6-000>.

³ Specifically, Bitcoin's average daily dollar volume in 2018 and 2019 was \$16.7BN and \$6.1BN (USD), respectively, as per historical data provided by CoinMarketCap.

B. Chowdhry (✉)

Indian School of Business, Hyderabad, India

e-mail: bhagwan@isb.edu

S. Kim

Santa Clara University, Santa Clara, CA, USA

e-mail: srkim@scu.edu

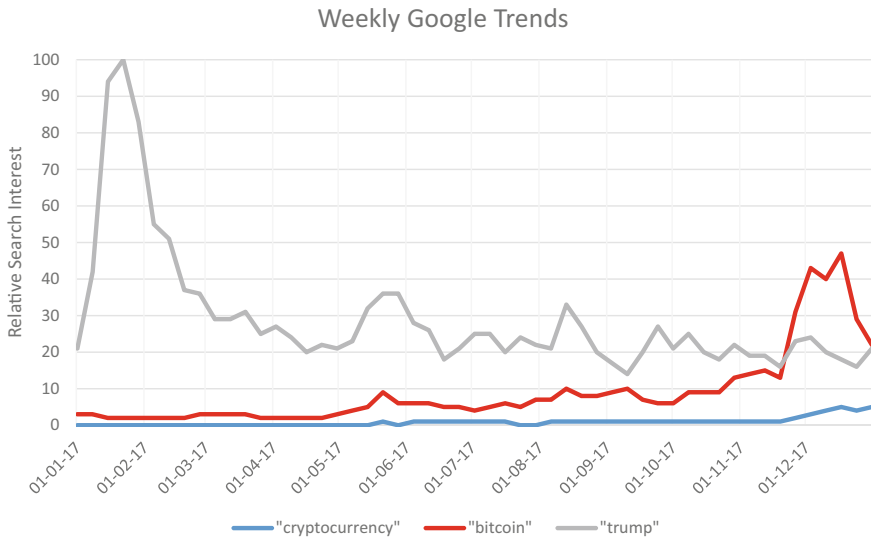


Fig. 1 Worldwide Interest in “Cryptocurrency” and “Bitcoin” (Source Google Trends). In this figure, we display the weekly, worldwide Google-search interest in the terms “cryptocurrency”, “bitcoin”, and “trump” over the period spanning January 1, 2017 through January 1, 2018. The relative search interest is scaled such that 100 represents the peak popularity for a term in the given time frame for provided search-term opportunity set: {“cryptocurrency”; “bitcoin”; “trump”}. See <https://trends.google.com> for further details

Despite the ongoing activity and interest in Bitcoin and the elusive “Satoshi Nakamoto,” whose whitepaper has now been translated into almost 30 different languages,^{4,5} this digital asset remains widely misunderstood. Although many now recognize Bitcoin as a disintermediated medium of exchange, there remains confusion surrounding who “runs” Bitcoin and how we can be confident in the security of a permissionless shared ledger in the way that we place confidence in Bank of America, as a trusted third party, to keep a proper record of funds. That is, Bitcoin was designed to provide “an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact without the need for a trusted third party” (Nakamoto 2008).⁶ But what exactly does this mean and how is it implemented?

Furthermore, the possibility of creating “trust” without using an intermediary such as a bank, a regulatory body, or government suggests many

⁴ The Github repository can be accessed on <https://github.com/wbnn/bitcoinwhitepaper>.

⁵ “Nakamoto” was even nominated for the 2016 Nobel Prize in Economics (Chowdhry 2016).

⁶ See page 1 of the original Bitcoin whitepaper, accessed on <https://bitcoin.org/bitcoin.pdf>.

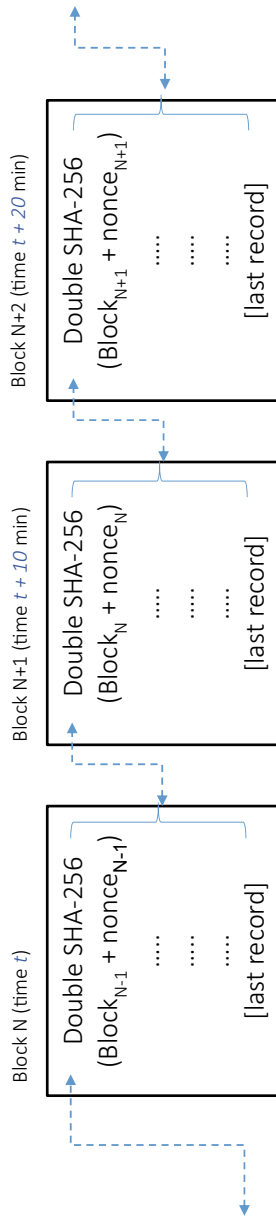


Fig. 2 The Bitcoin Blockchain. This figure graphically depicts a simplified version of the Bitcoin blockchain, which requires miners to find the winning “nonce” to close out a block of transactions records and to begin a new block. This nonce, when combined with the elements of the current block, must produce a double SHA-256 hash code with a minimum number of leading zeros. The first miner to find this nonce can close the current block and begin a new block, whose header will contain this nonce in addition to other pieces of information that link the new block to the prior one. Once this information is broadcast, other miners can easily verify that the new block was validly formed, and a new proof-of-work mining race begins

interesting use cases beyond financial transactions. This trustless security is the promise of distributed ledger technology in general. A public blockchain is a mechanism by which to store data, distributed across vast, peer-to-peer networks. Transactions or records are grouped into blocks that are chained to each other using cryptographic links. Blocks are verified and synchronized across many nodes only through agreement across various nodes of the network, a mechanism known as distributed consensus. Blockchain delivers immutability (i.e., it is impossible to modify past blocks of data unless a majority of nodes in the network collude, which is an unlikely and expensive situation as the network scales), and a reliable provenance of the transaction paths.

The focus of this chapter is to provide a layman's guide to Bitcoin and to explain the basic underpinnings of the consensus protocol (also known as the Nakamoto Consensus) that secures transactions on the Bitcoin blockchain.^{7,8}

1 An Intuitive Introduction

1.1 Numbers

Numbers are everywhere. We use numbers to count. Humans, because they have ten fingers, started using a decimal system that requires ten symbols: 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9. When we need to represent a number larger than 9, we begin using a second digit (i.e., 10, 11, 12 to 99) and then proceed to a third digit (i.e., 100, 101, 102) and so on.

Computers employ a binary system that uses only two symbols: 0 and 1. The counting in a binary system proceeds as 0, 1, 10, 11, 100, 101, and so on. Thus, 10 in binary is equivalent to 2 in the decimal system.

The binary system has a big advantage in that it is easy to make hardware that can be represented by 0s and 1s (e.g., OFF and ON). So, it is easy to build computers that only require switches that can be turned off and on. Furthermore, 0 and 1 can also represent False and True, or No and Yes, allowing us to represent any logical statement as a sequence of 0s and 1s. This was the remarkable insight by one of the greatest mathematicians, Claude Shannon, who founded information theory in the 1930s and 1940s.

⁷ See Kim et al. (2018) for a generalized overview of cryptocurrencies, including but not limited to Bitcoin.

⁸ For examples of other blockchain-based platforms and additional use cases, "Blockchain for Dummies, 2nd Edition" (Laurence 2019).

Alan Turing around the same time also showed that any computer language could be reduced to a sequence of 0s and 1s.

In the late 1950s, it was discovered that the basic building block of life itself only required four molecules (A, T, G, C) that makes the DNA molecule, and thus, the genome of any living being could also be represented by a number with four symbols, which could be 0 (A), 1 (T), 2 (G), and 3 (C). Of course, any number in a system of base 4 could equivalently be represented as a binary number. It is easy to see that any language with a finite set of letters in the alphabet could also be equivalently represented as a number in a binary system. For example, the English language would require 52 symbols to represent all 26 letters of the alphabet in lower and upper case, plus a few additional symbols to denote punctuation marks such as the period and the comma. In fact, our experience with fossil records has indicated that DNA can be preserved for hundreds of thousands of years without requiring any energy sources. Recent exciting advances are, in fact, exploring the use of information preservation using DNA (Lee 2019).⁹

So then, a binary sequence could represent a (i) number, (ii) any message in any language or any document that could be represented as a digital file, (iii) a person identified by her genomic sequence, or (iv) any set of logical instructions—e.g., if X then Y else Z—or even any detailed algorithm or a computer program.

In order to economize on the number of digits, a hexadecimal system is often used instead of a binary system. Since a hexadecimal system requires sixteen symbols, “0” to “9” plus “a” to “f” are typically used. So “e3b,” for example, is a number in the hexadecimal system which is equivalent to $(14 \times 16^2) + (3 \times 16^1) + (11 \times 16^0) = 3643$ in the decimal system and 111000111011 in the binary system.

1.2 A Unique “Fingerprint” Number for Everything Interesting (SHA256 Hash)

We have now established that almost every interesting thing that exists (e.g., a document, a transaction, a computer program, a person, a book, the first edition of Encyclopedia Britannica, a movie, all of Wikipedia at the current moment, and so on) can be represented as a number. Suppose that each one of these things could be assigned an ID or a unique number (i.e., a *fingerprint*). How many binary digits would be sufficient to represent all such objects? The

⁹ See, for instance, <https://www.scientificamerican.com/article/dna-data-storage-is-closer-than-you-think/>.

number of interesting objects for which we need IDs might be a very large number, say a trillion x trillion x trillion (i.e., $10^{12} \times 10^{12} \times 10^{12}$, since 10^{12} in the decimal system is one trillion). This sequence can also be written as $(10 \times 10 \times 10)^{12}$. Since $2^{10} = 1024$, which is approximately equal to $10 \times 10 \times 10$, a trillion x trillion x trillion can also be expressed as $(2^{10})^{12}$ or 2^{120} . That is, 120 binary digits would be enough to represent all IDs required to uniquely identify one trillion x trillion x trillion objects.

Suppose that, rather than using just 120 digits, we allow each ID to be represented by a 256-digit binary number. This extension allows an incredibly large number of IDs, approximately equal to the number of atoms in one trillion solar systems. Surely, if we were to distribute one trillion x trillion x trillion IDs using a 256-digit binary number, the number of IDs issued would be a vanishingly small fraction of the total potential IDs available. Furthermore, if IDs were issued randomly from the available pool of 2^{256} IDs, knowing the ID would provide no indication as to which object (which itself is represented by a number) it represents. The SHA-256 hash, where SHA refers to Secure Hash Algorithm, provides such an ID.

For each object, the SHA-256 hash code is fixed. That is, hash functions are deterministic, and the same input always yields the same hash code. However, because there are 2^{256} possible SHA-256 codes, it would be incredibly difficult to guess which object(s) generated a given SHA-256. Furthermore, given the complex nonlinear nature of the SHA-256 hash function, two seemingly similar inputs can yield SHA-256 hash codes that are completely different, which further complicates the guesswork in finding the original object that matches a provided SHA-256 hash code.

In cryptography, numbers are often expressed in a hexadecimal system, which requires sixteen symbols. Thus, a SHA-256 hash that requires 256 binary digits can be represented in a hexadecimal system using only 64 digits, where each digit is represented by one of the numbers 0 to 9 or letters “a” to “f.”

For example, the SHA-256 hash¹⁰ of the message “Prof. Bhagwan Chowdhry can explain Blockchain” is

8f8d3f14177bf664e9748cc790d739a71a7d06b2907844a309b43478e45c7124

If the original message, instead, were “Prof Bhagwan Chowdhry can explain Blockchain” its SHA256 hash is

db5bebb9f82d1040e96486b8573d91f7e205a025f41b8c7733dfa229c649b121

¹⁰ A SHA-256 hash calculator can be accessed on <https://andersbrownworth.com/blockchain/hash>.

Notice that the two hash codes are completely different even though the only difference between the first and second messages is that the word “Prof” is followed by a period in the former but not in the latter. Similarly, if the message were “Prof Bhagwan Chowdhry can explain blockchain” (where the only difference now is that the word “blockchain” is spelt with a lower case “b”), the SHA-256 hash is again completely different from the two hash codes shown above:

```
127b89ac962afc11e4c79fa265789633de209abac7400fd220c1268462ebde6c
```

Thus, if some original document's SHA-256 hash were presented alongside a tampered version of the original document, the SHA-256 hash of the tampered document would not match the hash code of the original document, signaling immediately that the document has been altered. This is the fundamental mathematical insight that many blockchain-based use cases exploit, whereby the authenticity of documents, such as birth certificates, transcripts, records of property ownership, etc., must be proven quickly (sometimes, nearly instantly) and efficiently without relying on cumbersome and expensive notary services, which often take a long time.

Hash codes are also referred to as checksums, since they allow you to prove access to or knowledge of the contents of a particular document without revealing any information contained in the document. That is, because it is nearly impossible to produce the correct SHA-256 hash without knowing the contents of the document, producing the SHA-256 hash code provides credible evidence that one indeed has the document in question.

For instance, consider the following tweet posted by Julian Assange of Wikileaks on October 16, 2016:

pre-commitment 1: John Kerry

```
4bb96075acadc3d80b5ac872874c3037a386f4f595fe99e687439aab0219809
```

Specifically, Julian Assange was sending a message to John Kerry (the, then, Secretary of State to President Barack Obama of the United States) to convey that he had successfully hacked into the State Department's documents, the proof of which was the SHA-256 hash. After seeing this message, Secretary Kerry could easily compute the SHA-256 hash code himself and would be able to confirm whether Julian Assange has access to the secret document in question.

Furthermore, the Twitter message is public for anyone to see and yet, no contents of the secret document are revealed by posting the SHA-256

hash. Thus, the SHA-256 hash code provides what is termed a *zero-knowledge proof*.

1.3 Digital Signatures Using Cryptography

There is often a need to transmit complete contents of messages as opposed to just its SHA-256 hash, from which the original message cannot be recovered. To do so electronically, without the contents of the message being intercepted by a third party, requires that the message be encrypted. Encryption refers to a two-way process whereby the original message is transformed using some one-to-one deterministic algorithm. For example, “I love you, Chris” could be encrypted as “J mpwf zpv Disjt” using a simple substitution in which each letter of the original message is substituted by the next letter in the alphabet. Anyone intercepting the message would either have to guess or know how the original message was transformed (i.e., the encryption key). An attempt to transmit the key electronically runs the risk that not only the message but also the key may be intercepted by a rogue third party.

A solution to this problem was proposed by Diffie and Hellman and further refined by Merkle in the mid-1970s (Levy 2002).¹¹ The basic idea is to generate a pair of keys, termed Public and Private keys. One of the keys is used to encrypt a message, and the other key is used to decrypt it back to its original form. For example, in order to make sure that the message “I love you, Chris” is not intercepted by a rogue third party, such as a spouse, Chris generates a pair of Public and Private keys. Chris publishes the Public key, making it available to everyone (including her secret lover), but keeps the Private key *private* by not revealing it to anyone (including her secret lover). Using Chris’s Public key, the message “I love you, Chris” is transformed into unintelligible gibberish which can be converted back to the original message only by using Chris’s corresponding private key, which only Chris knows. Thus, we avoid the problem of transmitting the key required for decryption.

However, another potential problem is that Chris cannot be sure who actually sent the encrypted message because the Public key is available to everyone. For example, the unhappy spouse could transmit a message, “Never talk to me again, Chris.” Although only Chris would be able to decrypt the message, she would not know whether the message was sent by her lover or the unhappy spouse, who is trying to sabotage the illicit love affair.

¹¹ See, for instance, “Crypto: How the Code Rebels Beat the Government Saving Privacy in the Digital Age” (Levy 2002).

To get around this issue, we can encrypt the message twice, first with Chris's Public key and then again with a different Private key, which also has a corresponding Public key that can be published publicly. Thus, the message has now been encrypted twice. Chris and her secret lover keep their respective Private keys private, while publicly publishing their respective Public keys.

When Chris receives a twice encrypted message, first she decrypts it using her lover's Public key—which is public—and then again using her own Private key. The unhappy spouse can no longer meddle because he does not know the Private key with which to first encrypt the original message.

In essence, any message can be digitally signed by encrypting the message with one's Private key. The encrypted message itself, which can only be decrypted by the corresponding Public key, can confirm that it must have been sent only by the person who knows the Private key. For example, one could prepend the message with the statement: "This message has been signed by Your Name." By adding this sentence to the beginning of the secret message "I love you, Chris" and, in addition, announcing publicly that all valid messages will begin with "This message has been signed by Your Name," Chris can easily verify the provenance of the message.

2 Bitcoin: A Bird's-Eye View

Bitcoin transactions are recorded and secured on a public, blockchain-based distributed ledger.¹² In this system, transaction records in the ledger are grouped into blocks, whereby a new block is formed approximately every ten minutes and is cryptographically linked to the prior block using hashed information from the prior block. The ledger is then replicated and maintained across numerous participants and systems, referred to as nodes, who do not need permission from a central authority to access these records.¹³ In contrast, random users cannot gain read/write access to Bank of America's ledger without permission. With a sufficient number of distinct nodes in the Bitcoin network, a dishonest node's attempts to alter transaction records or to validate faulty transactions will be overridden by the majority consensus of honest nodes.

Given the lack of central leadership in such networks, where participants can freely enter and exit in a permissionless and leaderless fashion,

¹² See Kim and Sarin (2018) for a generalized overview of distributed ledger technology, including but not limited to blockchain-based ledgers.

¹³ In fact, we (the authors) maintain a node on the Ethereum network, which is another public, blockchain-based distributed ledger!

these systems are also often called *decentralized autonomous organizations* (i.e., DAOs). However, to date, Bitcoin is the only cryptocurrency that comes closest to being a true DAO. That is, not only is participation in the Bitcoin network leaderless and permissionless (i.e., anyone can choose to maintain a replicated copy of the ledger or even choose to participate as a validator/miner who can add new blocks of records to the ledger), the ongoing upkeep of the underlying protocol is also managed in a leaderless and permissionless fashion. Specifically, anyone can submit Bitcoin Improvement Proposals (i.e., BIPs) to propose updates or changes to the network protocol, which will be implemented in the mining software once it has been accepted by the community of active Bitcoin miners.¹⁴

These updates to the network protocol, upon agreement by active miners, are known as *soft forks*. For instance, one proposal (BIP-0098)¹⁵ led to the adoption of a more efficient Merkle hash tree to summarize transactions in each new block header. Another proposal (BIP-0065)¹⁶ led to the implementation of an additional security feature to specify a time lock on transactions. Not all BIPs reach consensus across miners, which can lead to what is known as a *hard fork*, whereby a new cryptocurrency is created with the desired feature that was not accepted by the Bitcoin mining community as a whole. For instance, BIPs to adjust block sizes in the Bitcoin blockchain were not generally accepted by Bitcoin miners, which ultimately led to the creation of Bitcoin Cash (BCH),¹⁷ the first of many hard forks on Bitcoin.

In contrast, updates to Ethereum's network protocol are ultimately decided upon by a consolidated team of developers, known as *Ethereum Core Devs*. Although anyone in the general community can submit an Ethereum Improvement Proposal (EIP), the decision to implement the proposed updates does not rest on the general consensus of the Ethereum mining community. Thus, although participation as a node on the Ethereum network is leaderless and permissionless, the upkeep of the underlying protocol that determines the rules of the network is not. Other cryptocurrencies are even more centralized in the upkeep of their underlying protocol, lacking the open discussions and invitation of improvement proposals from their respective communities.

¹⁴ For instance, see BIP-0001, accessed on <https://github.com/bitcoin/bips/blob/master/bip-0001.mediawiki>, which established the concept and guidelines for all subsequent BIPs.

¹⁵ Accessed on <https://github.com/bitcoin/bips/blob/master/bip-0098.mediawiki>.

¹⁶ Accessed on <https://github.com/bitcoin/bips/blob/master/bip-0065.mediawiki>.

¹⁷ See <https://www.bitcoincash.org/>.

3 The Bitcoin Blockchain

The Bitcoin blockchain is secured by a hashcash-based proof-of-work (PoW) protocol,¹⁸ which requires solving a computationally difficult puzzle to close out the current block of records and begin a new one.

Hashing is a critical part of the Bitcoin protocol. Hashing is used in the generation of *private keys*, which are used as part of a *digital signature* to verify a user's legitimate access to Bitcoin funds. It is also an important part of the proof-of-work puzzle that miners must solve to close out one block and begin another.

Specifically, the Bitcoin blockchain uses the double SHA-256 hash function (also known as SHA-256²). As we discussed earlier in Sect. 1.2, the SHA-256 function is a *secure hash algorithm* that generates a 256-bit hash code (also known as the checksum), which is typically expressed as a 64-digit number in base 16. Double SHA-256 repeats the hash process by entering the SHA-256 hash code from the first iteration into the SHA-256 hash function once more, thereby producing a double SHA-256 hash code. Inputs to the SHA-256 hash function can be of any length, but the resulting hash code is always a 256-bit number.

Overall, the safety of the Bitcoin network, which allows anyone read/write access to its distributed ledger of transaction records, hinges critically on forcing validators in the network to solve computationally taxing cryptographic puzzles that would be nearly impossible for any single node (or colluding group of nodes) to resolve in a reasonable timeframe. We now proceed to explain the role of Bitcoin miners in validating transactions and maintaining the integrity of this permissionless, trustless system.

3.1 Nonces and Miners

Miners work to validate transactions requests, which are checked against past transaction records on the Bitcoin blockchain. Once verified, valid transactions wait in a memory pool (a.k.a., *mempool*) until they are added to a block that has been closed by a miner and confirmed by the network.

To successfully close a block and begin a new one, miners search for an arbitrary value called a *nonce*, that, when combined with the elements of the current block, must produce a double SHA-256 hash code with a minimum

¹⁸ See, for instance, "Hashcash—A Denial of Service Counter-Measure" (Back 2002), accessed on <http://www.hashcash.org/papers/hashcash.pdf>.

number of leading zeros. Thus, this nonce is difficult to find, but simple to verify once a solution is offered.

The first miner to find a winning nonce can close the current block, and “mine” a new block, whose header will contain this nonce in addition to other pieces of information that cryptographically chain the new block to the prior one.¹⁹ The header also contains code to generate a *block reward* to compensate the miner for forging this new block.²⁰ Once this new block is broadcast to the network, other miners can easily verify that the nonce is valid, and a new proof-of-work mining race begins. The Bitcoin protocol is designed to dynamically adjust the difficulty level based on rolling average block times such that a new block is formed approximately every ten minutes.

The design choice in setting a ten-minute block time is specific, though not limited, to Bitcoin. For instance, Ethereum block times typically average around 15s.²¹ The block-time choice entails a tradeoff between achieving a faster first confirmation (i.e., the time required for a pending transaction to first be included in a valid block) and dealing with the ensuing chain splits, whether accidental or intentional. That is, faster block times allow pending transactions to be added to the blockchain more quickly. However, since the proof-of-work puzzle is less computationally taxing, there is a greater likelihood that multiple miners will find winning nonces close to simultaneously, thereby causing a temporary split in the chain which must be resolved.

Overall, the inherent difficulty in solving for the proof-of-work nonce we described above is what secures the integrity of this trustless, permissionless ledger, since a group of malicious nodes is unlikely to have the computing power to resolve a series of nonces to successfully alter transaction records that are multiple blocks deep. Thus, a dishonest node’s attempts to validate faulty transactions or to alter the contents of a prior block would be overridden by the majority consensus of properly functioning nodes.

3.2 Network Congestion and Wait Times

Although Bitcoin blocks, on average, close every ten minutes, transactions in a closed block are typically not considered “confirmed” for about an hour after being added. The reason is that, for added security measures, Bitcoin

¹⁹ Figure 2 provides a graphical representation of this process. Additional visual representations can be accessed on <https://andersbrownworth.com/blockchain/>.

²⁰ The current block reward, as of February 5, 2020, is 12.5 BTC, and is estimated to reduce to 6.25 BTC by May 2020. See “Countdown to the Bitcoin Halving”, which can be accessed on <https://bravenewcoin.com/insights/countdown-to-the-bitcoin-halving>.

²¹ See <https://etherscan.io/chart/blocktime>.

clients (i.e., end-user software to facilitate sending and receiving Bitcoin) typically set their thresholds such that a transaction is not officially “confirmed” until it is six blocks deep.²²

However, during times of high network volume, transactions may take hours to be added and officially confirmed. Specifically, the Bitcoin protocol limits block sizes to one megabyte. As Bitcoin usage has increased, average block sizes have grown dramatically over time, and in recent years, we observe that blocks often reach their maximum size limit.²³ If the collective size of pending transactions in the memory pool exceeds the one megabyte block-size limit, then unchosen transactions must inevitably wait to be added to a subsequent block. Based on an average transaction size of 570 bytes (Moos 2019),²⁴ an influx of 100,000 network messages would result in wait times in excess of five hours to clear and confirm all transactions in the memory pool.

Because miners also collect transaction fees attached to each transaction record, they are incentivized to gather the transactions with the highest fee-to-size ratios. Bitcoin clients tend to dynamically adjust fees for end-users based on the size of the transaction and extant network conditions. They also allow users the option to pay an enhanced priority fee, which makes a miner more likely to select that particular transaction from the memory pool when forming the latest block. Thus, pending transactions with lower fees risk sitting in the memory pool for long periods of time when network volume is high.

Finally, we note the vast differences in speed and confirmation times when transacting via a peer-to-peer Bitcoin client versus on a crypto-exchange. Because of the frustration and intractability inherent in maintaining an evolving limit order book when traders are forced to wait in excess of an hour to know whether they have the funds to bid or the assets to offer, exchanges simply maintain a centralized ledger of so-called *off-chain* transactions (as distinct from *on-chain* transactions). That is, transactions to move BTC to and from an exchange are recorded on the Bitcoin blockchain, and thus, are not instantaneously confirmed. However, transactions *within* an exchange

²² As of November 8, 2019, Coinbase has lowered this threshold to three blocks, which suggests that many Bitcoin clients may follow suit. See <https://blog.coinbase.com/announcing-new-confirmation-requirements-4a5504ba8d81>. In comparison, Coinbase requires an Ethereum transaction to be 35 blocks deep before considering it confirmed. See <https://help.coinbase.com/en/coinbase/trading-and-funding/sending-or-receiving-cryptocurrency/why-is-my-transaction-pending.html>.

²³ See, for instance, <https://bitinfocharts.com/comparison/bitcoin-size.html>.

²⁴ See <https://cryptoslate.com/bitcoin-transactions-per-block-at-all-time-highs/> for a discussion on average transaction sizes and average number of transactions per block.

occur off chain, and can be confirmed swiftly since the exchange's permissioned ledger does not require a computationally taxing proof-of-work puzzle to be solved to add a transaction.

Overall, scalability remains a challenge to the widespread adoption of Bitcoin as a medium of exchange. In the following section, we describe other design choices intended to increase the throughput of transactions while maintaining the integrity of transactions records.

4 Ledger Design

Introduced by Bitcoin and popularized by its success, the most widely used choice in public-ledger design is currently a blockchain-based ledger predicated on a proof-of-work consensus protocol. However, there are other types of ledgers and consensus protocols, which differ in their efficacy and propriety based on whether the ledger is intended to be public or private. We now proceed to a brief overview of these various design choices.

4.1 Which Consensus Protocol?

A key issue in a distributed record-keeping system is how to reach consensus across nodes without halting the system in the presence of a few faulty or malicious nodes. This property, known as Byzantine fault tolerance (BFT),²⁵ strives to achieve a balance between what is known as *liveness* (i.e., allowing transactions to occur) and *safety* (i.e., preventing faulty transactions from occurring). A proof-of-work based consensus mechanism is one commonly used choice in designing a Byzantine fault-tolerant network.

As we discussed, the Bitcoin blockchain is secured by a proof-of-work consensus protocol, whereby miners work to solve a computationally intense puzzle to add new blocks of transaction records to the existing chain. Proof of work, by design, is a slow and laborious process, and more recent algorithms are moving to alternative consensus mechanisms in an attempt to achieve greater scalability in light of the bottlenecks occurring during times of high network volume.

For instance, the proof-of-stake (PoS) consensus protocol has been a popular alternative, whereby a node or subgroup of nodes is selected to validate the next set of transactions to be added to the ledger. That is, rather

²⁵ See Lamport et al. (1982) for a discussion of the Byzantine generals problem and its applications to reliable distributed computing.

than having an entire pool of miners racing to solve a computationally taxing puzzle to win the right to add transactions to the existing ledger, a proof-of-stake system selects the next validator(s), typically with a randomized component, based on some relative stake in the system. For instance, the stake may be measured by the sheer size of a node's stake (i.e., wealth in native tokens),²⁶ or by a combination of the node's stake and age of stake (i.e., how long the node has held these native tokens).²⁷

Some modifications allow for lightweight nodes (i.e., nodes that don't maintain full copies of the blockchain database) with relatively little stake to lease their stake to full nodes on the network (known as *leased proof of stake*),²⁸ or use their stake to vote for the delegates to represent them (known as *delegated proof of stake*).²⁹ Others have begun to implement punitive elements to their proof-of-stake protocol (known as *punitive proof of stake*), whereby validators are not only rewarded for producing valid blocks but also punished for producing invalid ones.³⁰

4.2 Public Versus Private Ledgers

Our discussions on various considerations in ledger design and consensus protocol were predominantly fashioned with public (permissionless) ledgers in mind. However, with the rising popularity of Bitcoin, it has become in vogue to seek out blockchain solutions in a variety of settings, and firms have begun to consider private blockchain-based ledgers as well as other types of shared ledger designs for use within a group of permissioned entities.

For instance, Hyperledger,³¹ started by the Linux Foundation in 2015, provides private blockchain solutions that do not require a native token to operate. This should allay the fears of many regulators, such as those in India, who welcome the potential of blockchain technology but want to discourage

²⁶ See, for instance, the BlackCoin whitepaper (Vasin, n.d.), accessed on <https://blackcoin.org/blackcoin-pos-protocol-v2-whitepaper.pdf>.

²⁷ See, for instance, the Peercoin whitepaper (King and Nadal 2012), accessed on <https://www.peercoin.net/whitepapers/peercoin-paper.pdf>.

²⁸ See, for instance, the WAVES whitepaper (2016), accessed on <https://medium.com/wavesprotocol/waves-whitepaper-164dd6ca6a23>.

²⁹ See, for instance, the Steem whitepaper (2018), accessed on <https://steem.com/steem-whitepaper.pdf>.

³⁰ For an early discussion of a punitive proof-of-stake protocol, see *Slasher: A Punitive Proof-of-Stake Algorithm* Buterin (2014), accessed on <https://blog.ethereum.org/2014/01/15/slasher-a-punitive-proof-of-stake-algorithm/>.

³¹ <https://www.hyperledger.org/>.

the proliferation and use of cryptocurrencies. Depending on the permission settings and selected consensus mechanism, the corresponding network may be categorized as public (e.g., Bitcoin) or private (e.g., one built on Hyperledger Fabric).

Naturally, some features used in distributed ledger design can be implemented in a practical and sensible manner for use in a private ledger. For instance, the simple act of grouping transactions into blocks, which are chained in sequence, is not a particularly novel idea and has been implemented by many in their own private ledgers. For instance, in our finance records, we often create blocks (perhaps one for each year). Then, at the end of the calendar year, we close out the 2019 block and form a new 2020 block whereby the first element is chained to the last element of the prior block, thereby forming the simplest of blockchains.

However, some design features used in many public distributed ledgers, such as the proof-of-work consensus protocol, are clearly inappropriate in a private ledger design.

4.3 Blockchain-Based Ledgers Versus Directed Acyclic Graphs

Although, blockchain-based ledgers are still currently the most widespread design choice, not all distributed ledgers are blockchain-based, and more recently, other implementations are being explored in an attempt to achieve greater scalability. Specifically, a blockchain-based design requires synchronous consensus, such that blocks are agreed upon and added linearly in a chronological fashion. In an effort to mitigate bottlenecks inherent in such a design, some recent projects are moving away from a blockchain-based structure and toward a design that allows for asynchronous agreement to validate pending transactions.

For instance, one structure makes use of a Unique Node List,³² whereby overlapping subsets of nodes asynchronously reach consensus until the entire network reaches agreement. Another such structure makes use of a directed acyclic graph (DAG),³³ whereby each transaction must select other transactions to validate. A pending transaction is ultimately confirmed as it is repeatedly selected for verification and is nestled more deeply in the DAG.

³² See, for instance, Chase and MacBrough (2018) for an analysis of Ripple's XRP protocol, accessed on <https://arxiv.org/pdf/1802.07242.pdf>.

³³ See, for instance, the Byteball whitepaper (Churyumov, year unknown), accessed on <https://obyte.org/Byteball.pdf>.

Some nascent projects have also been researching the ability of this DAG-based structure to handle the throughput required of an IoT (Internet of Things) network, which would be impossible to handle with current blockchain-based designs.³⁴

5 Concluding Remarks

Now, more than ten years following its arrival, “Bitcoin” has quickly become a natural part of colloquial speech. Its widespread popularity has also brought to prominence terms such as “blockchain” and “distributed ledgers.” But the rapid adoption of these terms by the general public has also been fraught with many misunderstandings and ill-conceived use cases. Our hope, with this chapter, is to provide a layman’s guide to Bitcoin and to shed light on the basic mechanisms underlying the Bitcoin blockchain and other more general considerations in distributed ledger design.

References

- Back, Adam. 2002. *Hashcash—A Denial of service counter-measure*, August 1. Accessed on <http://www.hashcash.org/papers/hashcash.pdf>.
- BraveNewCoin.com. 2020. *Countdown to the Bitcoin Halving*, February 5. Accessed on <https://bravenewcoin.com/insights/countdown-to-the-bitcoin-halving>.
- Buterin, Vitalik. 2014. *Slasher: A punitive proof-of-stake algorithm*, January 15. Accessed on <https://blog.ethereum.org/2014/01/15/slasher-a-punitive-proof-of-stake-algorithm/>.
- Chase, Brad, and Ethan MacBrough. 2018. *Analysis of the XRP ledger consensus protocol*, February 21. Accessed on <https://arxiv.org/pdf/1802.07242.pdf>.
- Chowdhry, Bhagwan. 2016. *I (Shall Happily) accept the 2016 nobel prize in economics on behalf of Satoshi Nakamoto*, November 6. Accessed on https://www.huffpost.com/entry/i-shall-happily-accept-th_b_8462028.
- Churyuomv, Anton. *Byteball: A decentralized system for storage and transfer of value*. Accessed on <https://obyte.org/Byteball.pdf>.
- Coinbase. 2019. *Announcing new confirmation requirements*, November 8. Accessed on <https://blog.coinbase.com/announcing-new-confirmation-requirements-4a5504ba8d81>.

³⁴ See, for instance, an earlier IOTA whitepaper (Popov 2018), accessed on <https://www.iota.org/foundation/research-papers>.

- Friedenback, Mark, Kalle Alm, and BtcDrak. 2017. *BIP 98: Fast Merkle trees*, August 24. Accessed on <https://github.com/bitcoin/bips/blob/master/bip-0098.mediawiki>.
- Kim, Seoyoung, and Atulya Sarin. 2018. Distributed ledger and blockchain technology: Framework and use cases. *Journal of Investment Management* 16 (3): 90–101.
- Kim, Seoyoung, Atulya Sarin, and Daljeet Viridi. 2018. Crypto-assets unencrypted. *Journal of Investment Management* 16 (2): 1–31.
- King, Sunny, and Scott Nadal. 2012. *PPCoin: Peer-to-Peer crypto-currency with proof-of-stake*, August 19. Accessed on <https://www.peercoin.net/whitepapers/peercoin-paper.pdf>.
- Lamport, Leslie, Robert Shostak, and Marshall Pease. 1982. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems* 4 (3): 382–401.
- Laurence, Tiana. 2019. *Blockchain for dummies*, 2nd edn, May 7.
- Lee, Sang Yup. 2019. *DNA data storage is closer than you think*, July 1. Accessed on <https://www.scientificamerican.com/article/dna-data-storage-is-closer-than-you-think/>.
- Levy, Steven. 2002. *Crypto: How the code rebels beat the government saving privacy in the digital age*, January 15.
- Moos, Mitchell. 2019. *Bitcoin transactions per block at all-time highs*, April 8. Accessed on <https://cryptoslate.com/bitcoin-transactions-per-block-at-all-time-highs/>.
- Nakamoto, Satoshi. 2008. *Bitcoin: Peer-to-peer electronic cash system*, October 31. Accessed on <https://bitcoin.org/bitcoin.pdf>.
- Popov, Serguei. 2018. *The Tangle*, April 30. Accessed on <https://www.iota.org/foundation/research-papers>.
- Steem. 2018. *An incentivized, blockchain-based, public content platform*, June. Accessed on <https://steem.com/steem-whitepaper.pdf>.
- Taaki, Amir. 2011. *BIP 1: BIP purpose and guidelines*, August 19. Accessed on <https://github.com/bitcoin/bips/blob/master/bip-0001.mediawiki>.
- Todd, Peter. 2014. *BIP 65: OP_CHECKLOCKTIMEVERIFY*, October 1. Accessed on <https://github.com/bitcoin/bips/blob/master/bip-0065.mediawiki>.
- Vasin, Pavel. n.d. *BlackCoin's proof-of-stake protocol V2*. Accessed on <https://blackcoin.org/blackcoin-pos-protocol-v2-whitepaper.pdf>.
- Waves Protocol. 2016. *WAVES Whitepaper*, April 1. Accessed on <https://medium.com/wavesprotocol/waves-whitepaper-164dd6ca6a23>.
- Zmudzinski, Adrian. 2019. *Bitcoin ATMs worldwide hit new milestone, surpassing 6,000*, November 17. Accessed on <https://cointelegraph.com/news/bitcoin-atms-worldwide-hit-new-milestone-surpassing-6-000>.



An Introduction to Cryptocurrencies

Jeongmin Lee, Christine A. Parlour, and Uday Rajan

Cryptocurrencies typically refer to the instruments issued by private parties on a “blockchain,” a type of distributed ledger used for record-keeping. When used in exchange for goods and services, cryptocurrencies act as a means of payments similar to credit cards, which were introduced in the United States in the 1970s. The key difference between traditional digital payment methods and cryptocurrencies is that blockchain technology allows anyone to issue their own cryptocurrency, while credit cards are issued by well-known banks or institutions.¹ As a practical matter, cryptocurrencies (a.k.a. tokens or coins) are digitally native (i.e., exist solely in digital form) and settled outside existing institutional channels.

¹ The Erc-20 standard on the Ethereum blockchain has facilitated token issuance.

J. Lee

Washington University, University of Arizona, St Louis, MO, USA

e-mail: jlee89@wustl.edu

C. A. Parlour (✉) · U. Rajan

University of California, Berkeley, CA, USA

e-mail: parlour@haas.berkeley.edu

U. Rajan

e-mail: urajan@umich.edu

Bitcoin, one of the first cryptocurrencies, came into existence on January 3, 2009, when the “genesis” (i.e., first) block was mined. Since then, consumers all around the world have been exchanging value with these instruments. Transactions have been verified and settled in agents such as Bitcoin “miners” participating in decentralized system. Today, there are over 2,000 actively traded cryptocurrencies and many others have been issued and failed. As of January 10, 2019, the total market cap of cryptocurrencies is over 120 billion US dollars, of which Bitcoin constitutes over 60 billion dollars.²

The rapid rise of these new instruments raises a number of questions. Is a cryptocurrency something economically new or is it merely new technological packaging of something familiar? If cryptocurrencies are something new, what economic effect will they have and what role will they play? What are the possible drawbacks? Is there a role for policy? The answers to these questions are fundamentally important for businesses and regulators. In this chapter, we attempt to synthesize the fast-emerging literature to provide an overview of the potential roles of and concerns related to cryptocurrencies. While none of our answers will be definitive, we make some suggestions for future research.

1 Potential Roles and Benefits

We begin by discussing the potential roles that cryptocurrencies can play and the possible benefits from them. First, cryptocurrencies can be viewed as a new type of currency or money, in that they can be used as a medium of exchange for various transactions. Cryptocurrencies then have value as long as they are widely accepted for transactions. Second, cryptocurrencies can be viewed as a new type of investible asset. Investors can participate in trading cryptocurrencies for hedging and speculative purposes. Third, cryptocurrencies can be used to raise funds for projects and funds. The market for initial coin offerings (ICOs) is growing rapidly and can present new opportunities for entrepreneurs and investors alike. Fourth, cryptocurrencies can be used to reduce transaction costs in making payments. They are being used for international remittances as a much cheaper alternative to the current banking system. Finally, we discuss implications of cryptocurrencies for central banking.

Payment Instruments

An effective transaction medium provides an efficient way to keep track of the pairwise relative values of every good and service (e.g., how many haircuts

² For an up-to-date list of cryptocurrencies and their market cap, see coinmarketcap.com, while deacoin.com provides a partial list of those that have failed.

or iPhones are equivalent to a lecture by a university professor). It also facilitates exchange between unknown or anonymous counterparties Narayana R Kocherlakota (1998). By convention, we expect money to be issued by sovereign nations, in the form of fiat currency.

Anything that is widely accepted for transactions can be used as money. Until the middle of the eighteenth century, playing cards were widely circulated in place of either specie or fiat currency in New France, Canada (now Québec). The colony was far from France, and having a medium of exchange was vital for small homesteads spread over the new territory to trade with one another. The card money was accepted and circulated for over a century.³ Few would accept an ace of spades for payment today. Throughout human history, means of payment has evolved, from gold and silver coins to paper currency to bank checks and credit cards. Viewed in this way, cryptocurrencies may be seen as the next step in the evolution of means of payment.

With Bitcoin being originally envisioned as a new payment instrument (Satoshi Nakamoto 2008), the first wave of literature casts Bitcoin in the context of currency models. For example, Susan et al. (2016) posit that Bitcoin has value because people use it for transactions, and that people believe it will be useful for transactions in the future. This latter part has two elements: First, consumers who accept Bitcoin now must believe that they can use it for their own needs in the future, and second, that the technology will not “fail.”⁴ In their framework the exchange rate (e.g., the dollar value of one Bitcoin) is determined by economic fundamentals, i.e., the supply and demand of Bitcoin. They also note that the existence of speculators with no use value will affect the pricing and adoption. The possibility of speculation (and hedging) makes Bitcoin (and other cryptocurrencies) a new investment opportunity.

A simple way to understand the “currency” framework is to recall the definition of money velocity. Suppose that q_t represents the quantity of real goods and services purchased in a period of time t , at a weighted average price of p_t . Then, by definition the total expenditure in that period is $p_t q_t$. This is true if the prices are quoted in USD, Japanese yen, or Bitcoin. We will denote the USD price as $p^{\$}$, and the Bitcoin price as p^{B} . Suppose that the total number

³ For more detail on this fascinating period of monetary history, see Roeliff Morton (1893) and Herbert Heaton (1928).

⁴ We discuss the literature on the possibility of technological failure in cryptocurrencies in the next section.

of mined Bitcoins is M_t^{B} . Then, the velocity of Bitcoin, V_t^{B} , is simply

$$V_t^{\text{B}} = \frac{p_t^{\text{B}} q_t}{M_t^{\text{B}}}.$$

We can rearrange this definition to get an understanding of how the exchange rate between Bitcoin and USD changes. Specifically,

$$\begin{aligned} M_t^{\text{B}} V_t^{\text{B}} &= p_t^{\text{B}} q_t \\ M_t^{\text{B}} V_t^{\text{B}} &= \frac{p_t^{\text{B}}}{p_t^{\text{S}}} (p_t^{\text{S}} q_t) \end{aligned}$$

To determine the exchange rate at time t , e_t , i.e., the USD price of 1 Bitcoin, we just use the fact that

$$p_t^{\text{B}} = \frac{p_t^{\text{S}}}{e_t}.$$

So, it immediately follows that

$$e_t = \frac{(p_t^{\text{S}} q_t)}{M_t^{\text{B}} V_t^{\text{B}}}.$$

The dollar value of one Bitcoin is thus determined by the dollar value of the total expenditure ($p_t^{\text{S}} q_t$) divided by the gross volume of Bitcoin ($M_t^{\text{B}} V_t^{\text{B}}$).

Investible Assets In spite of their name, there is considerable debate about whether to think of cryptocurrencies as currencies or assets. The distinction between these two instruments is crucial for regulatory purposes. In particular, whether a cryptocurrency is deemed a currency or a security determines whether it is subject to the SEC's securities regulations.⁵

Practically, the distinction between currencies and assets is somewhat imprecise. Although we usually think of a currency as a means of payment, from the point of view of an English investor, the Japanese Yen or US dollar is an investment with risky payoffs. Fiat money is usually considered to be a pure common value good with no fundamental value. From this perspective, the value of US dollars is due to a bubble, since it has no fundamental value. By contrast, assets either have an eventual use value (as in the case of claims on commodities) or are claims on cash flows that come from the sale

⁵ See the SEC's statement at <https://www.sec.gov/news/public-statement/digital-asset-securities-issuance-and-trading>.

of some underlying project or asset. Importantly, even though in equilibrium both fiat money and speculative assets have a value that everyone holds in common, with speculative assets different agents can have different information or endowment positions and hence can have different needs to trade and different willingnesses to pay. By contrast, most people are willing to pay one dollar for one dollar.

Wilko Bolt et al. (2016) provide a framework to value a cryptocurrency with both a use value as a payment instrument and a value as an asset, or a speculative value. Here, they start from the premise that there is a value for cryptocurrencies for making payments as in Susan et al. (2016). In addition they assume that the supply of the currency grows at a deterministic rate, as is the case with Bitcoin.

They observe that speculators may buy the cryptocurrency today in anticipation of future price increases. If enough speculators buy and hold the currency, then this effectively reduces the supply, and as is standard, supply reductions lead to increases in price. There is a feedback between what everyone expects the exchange rate to be, and the number of merchants and customers who adopt the currency, which feeds back into its value.

Recall the simple currency framework from earlier, where the dollar value of total expenditure is $p_t^{\$}q_t$ and the gross volume of Bitcoin exchanged is $M_t^{\text{B}}V_t^{\text{B}}$. The framework can be used to tie fundamentals to the Bitcoin exchange rate, providing micro-foundations for the quantity of Bitcoin, or the velocity.

For example, suppose that a proportion of the population buy and hold a cryptocurrency for speculative purposes, i.e., do not use it for transaction purposes. Call this amount of cryptocurrency Z_t^{B} . The velocity of these coins would be zero (it is not being used to buy or sell anything). Hence, one can decompose the average velocity of Bitcoin into a part that is actually used for payments and a part that is held for speculative purposes.

$$\tilde{V}_t^{\text{B}} = \frac{M_t^{\text{B}} - Z_t^{\text{B}}}{M_t^{\text{B}}} V_t^{\text{B}} + \frac{Z_t^{\text{B}}}{M_t^{\text{B}}} 0.$$

This immediately means that

$$e_t = \frac{\left(p_t^{\$}q_t\right)}{\left(M_t^{\text{B}} - Z_t^{\text{B}}\right) V_t^{\text{B}}}.$$

Clearly, the exchange rate (or the dollar price of Bitcoin) is higher in the dollar volume of payments that are made in the cryptocurrency, is smaller

the larger the supply of cryptocurrency is, and is higher the more currency is kept for speculative reasons. It is also lower when the velocity is higher. The velocity is determined in part by how many investors use the cryptocurrency, in other words, by network effects.

A payment system is a canonical two-sided market. Specifically, both parties to a transaction have to agree that they want to transfer value, and both parties have to agree on the medium to transfer that value. Further, the more people who agree on the ideal way to transfer value, the more effective a particular medium becomes. The latter is an example of a network effect, in this case a positive network effect. The more people who use a particular payment method, the more valuable it is.

For this reason, some researchers have focused on the importance of network effects in evaluating a cryptocurrency. Lin William Cong et al. (2018) view the world as one in which activities on the blockchain are valuable, and tokens or coins are claims on those activities. The equilibrium effects they characterize arise because an individual's benefit from blockchain activities is increasing in the total number of users (i.e., there is a positive network effect). Similar to Wilko Bolt et al. (2016), agents decide how many tokens they need for their activities, and also how many tokens they want to acquire because they anticipate future appreciation in value. Interestingly, more token appreciation leads to more token adoption, which lifts overall values. This framework neatly captures the relationship between token use and token value. Michael Sockin and Wei Xiong (2018) show that strategic complementarities in cryptocurrency demand can lead to either none or two equilibria, creating fragility.

Another way to determine if an asset adds value is to see if its returns are driven by any of the standard risk factors. Yukun Liu and Aleh Tsyvinski (2018) show that the returns of cryptocurrencies have no exposure to stock returns and common macroeconomic factors. Using a data set that comprises 222 cryptocurrencies with a market capitalization of over \$ 1 million on November 1, 2017, Albert Hu et al. (2018) consider what drives the return series. Table 1 reproduced from their paper shows the correlations with standard risk factors.

As is evident, on average cryptocurrency returns seem to be related to Bitcoin returns, and indeed, the authors indicate that the first principal component of the return series is highly correlated with Bitcoin returns. However, we emphasize that the quality of the data may not be good, especially in light of the documented price manipulation we allude to later.

In a mean-variance framework, the value of an asset depends in part on the diversification benefit that it adds to a portfolio. Alla Petukhina et al. (2018)

Table 1 Daily (top) and monthly (bottom) altcoin correlations with other assets

Correlation with	Mean	SD	Min	p25	p50	p75	Max
Bitcoin	0.174	0.113	-0.123	0.0926	0.177	0.251	0.645
Gold	0.0193	0.0474	-0.145	-0.0081	0.0146	0.0408	0.223
S&P 500 excess return	0.0045	0.0468	-0.224	-0.0184	0.0034	0.03	0.152
Observations	221						
Correlation with	Mean	SD	Min	p25	p50	p75	Max
Bitcoin	0.21	0.29	-0.382	0.0057	0.18	0.363	0.949
Gold	0.0611	0.186	-0.456	-0.0657	0.047	0.156	0.688
S&P 500 Excess Return	-0.0166	0.189	-0.671	-0.106	0.0179	0.0863	0.527
Observations	221						

This table shows summary statistics for altcoin correlation coefficients with Bitcoin, Gold, and the S&P500 Excess Return. There are 221 observations as Bitcoin is excluded. The top part of the table shows daily correlations, and the bottom part shows monthly correlations.

show that cryptocurrencies can improve the risk-return profile of some portfolios. We note, however, that from a portfolio construction point of view, the return distribution of cryptocurrencies is heavily skewed (i.e., is not normal).

Funding Methods

In an initial coin offerings (ICO), a new cryptocurrency is issued to raise funding for a project or firm. The ICO market is growing rapidly. Between 2014 and 2018, ICOs raised over 18 billion USD (Sabrina T. Howell et al. 2018).

Jeongmin Lee and Christine A. Parlour (2018) present a model to show the possible role of cryptocurrencies in this context. In that many ICOs were issued to fund projects or products, ICOs can be viewed as a form of crowdfunding, which allows consumers directly to participate in providing funding. Currently, venture capitalists and registered investors constitute the majority of the funding sources for small firms. The authors point out that when financial institutions make funding decisions, they do not take into account the consumer surplus that consumers enjoy after the products are produced. Hence, even absent other frictions, the institutions' funding choice faces an underinvestment problem, which prevents a project with positive total surplus but negative net cashflows from being undertaken.

Of course, after having committed to fund a project, consumers may decide that they no longer want to consume the good or use the service (in other words, they may suffer a negative "liquidity shock"). In this case,

having a market in which to sell their claims is valuable. Thus, a cryptocurrency can be viewed as a type of crowdfunding with a resale market. As such it is valued more highly by consumers and so leads to more efficient investment than funding by standalone banks or by traditional financiers that rely on a cash flow rule. Going further, the authors claim that some “frothiness” or speculation can be socially efficient, in that higher resale prices might induce more efficient investment, which could be socially beneficial.

An interesting implication of their framework is that empirical researchers should be cautious about interpreting the success (or lack thereof) of projects or ventures that are funded by cryptocurrencies. In particular, traditional cash flow measures that we typically use to assess whether a project or a venture is successful may not reflect the project choice that is socially optimal. One of the challenges for empirical work is to find a reasonable measure of the use value that some of the owners of cryptocurrencies may obtain from their ownership.

There is a fast growing literature on the economic effects of ICOs. Christian Catalini and Joshua S. Gans (2018) show that the ICOs generate value for entrepreneurs by generating competition among buyers. Jiasun Li and William Mann (2018) study how the network effect affects the future users of the cryptocurrencies’ platform. Jiri Chod and Evgeny Lyandres (2018) show that ICOs allow the entrepreneurs to transfer cash flow risk without diluting their control rights. Katya Malinova and Andreas Park (2018) study the optimal design of tokens in ICOs and compare them with equity. Similarly, Rodney Garratt and Maarten van Oordt (2019) characterize how ICOs change the incentives of entrepreneurs relative to debt or venture finance.

Reducing Transaction Costs

Although much of the attention of the world has focused on Bitcoin, it is useful to consider another system that has the capacity to supplant the international banking system, namely, the real-time gross settlement system Ripple and its associated cryptocurrency, the XRP.

Typically cross-border transfers operate through what is known as the correspondent banking system. If two banks do not have a direct relationship, they send value through an intermediary bank with whom they both do business (i.e., a correspondent bank). For banks on either side of the highly used corridors such as London to New York, the system is painless. However for those sending value outside these corridors (for example, Ottawa to Gaborone) the process is extremely expensive, as each intermediary requires compensation. In addition to the cost associated with moving value, there is also a cost associated with currency exchange, as the two currency pairs (Canadian dollars and the Botswana pula) are not easily convertible.

The Ripple solution is to introduce a cryptocurrency, the XRP, which can essentially be used to intermediate exchange between different currency pairs. The benefit of the cryptocurrency is that it is cheaper than the banking system for international remittances. The correspondent banking system described above is inefficient and expensive. To some extent, the value of this cryptocurrency is the reduction in transaction costs over the usual banking system.

Central Banking

In countries with high inflation (e.g., Venezuela, where the inflation rate was at one point over 25,000% last year), people are attracted to holding more stable foreign currencies like USD. When they cannot do so due to capital controls and other regulation, cryptocurrencies such as Bitcoin can be a possible alternative. In this sense, cryptocurrencies could theoretically compete with central bank's fiat money in troubled countries. Can a central bank issue its own cryptocurrency?

Max Raskin and David Yermack (2016) and references therein discuss the potential roles of cryptocurrencies and the blockchain technology behind them in central banking. One advantage of a cryptocurrency is that it can commit to limit the number of new coins that are issued. This can make the central bank cryptocurrency more appealing than the fiat money of the government and the central bank, which people no longer trust. Another interesting idea that is unlikely to happen in practice is that the central bank's cryptocurrency can be used to take deposits from retail customers directly. Replacing commercial banks entirely, it can thus eliminate the need for deposit insurance.

2 Limitations and Concerns

The blockchain technology that underlies cryptocurrencies has a few limitations. First, the possibility of technology failure limits the roles of cryptocurrencies. Second, transactions using cryptocurrencies are stochastic and subject to failure. Third, the market microstructure of cryptocurrencies is underdeveloped and unregulated. There are huge, persistent price discrepancies across different exchanges. Fourth, the cryptocurrency market is highly illiquid and subject to price manipulation. Finally, the surveys show that the real-world usage of cryptocurrencies is far from their theoretical promises, with a large proportion being used for illegal activity. Fifth, exchanges on which cryptocurrencies are traded are typically small and private, leading to particular risks of their own. Two examples include bitcoins being stolen when Mt. Gox

was hacked in 2014, and the CEO of Quadriga suddenly dying in 2019, with no one else being able to access account information.

Blockchain Failure The value of cryptocurrencies as payment instruments requires today's users to trust that the blockchain technology behind them will not fail. Although the technology is proposed as a safe and reliable way to keep records, it is far from being fool proof. In 2017, Bitcoin suffered two "forks;" i.e., situations that led to multiple versions of the distributed ledger. Naturally, records become unreliable when there are multiple versions of history. Bruno Biais et al. (2018) present a game-theoretic framework that highlights the importance of coordination among miners. They find that there can be multiple equilibria, and an equilibrium with a fork event can exist.

Eric Budish (2018) provides a simple framework to show the cost of the incentive constraint inherent in blockchain technology. The rewards for the miners must be large relative to their benefits from attacking it, so making the blockchain incentive compatible is very expensive. Furthermore, he shows that there is a limit to the role that Bitcoin can play in the economy, since it would be attacked if it becomes sufficiently important. Joseph Abadi and Markus Brunnermeier (2018) point out that no blockchain can achieve correctness, decentralization, and cost efficiency all at the same time. They compare blockchains with centralized record-keeping and show that the two systems can complement each other.

Peter Zimmerman (2018) carefully considers the joint impact of a blockchain's settlement capacity and Bitcoin's use as a means of payment. He argues that the existence of limited settlement capacity reduces the value of a cryptocurrency as a medium of exchange. Indeed competition for settlement from speculators leads to a lower value as a medium of exchange and higher price volatility. His work emphasizes that even though cryptocurrencies may exhibit some familiar economic characteristics, the novel aspects of the technology itself can lead to different economic outcomes.

Transaction Failure

We should point out that one of the benefits of physical "cash" is that it can be used for delivery against payment. That is, one can lay a dollar on the counter in exchange for the goods that one wants. With credit or debit cards, the act of payment also acts as a receipt of the transaction. The receipt can in turn be used for refunds or to contest a transaction. By contrast, with Bitcoin, because the verification is decentralized, payment is stochastic. You may hand over a good to someone in exchange for a transfer to your wallet, receive a confirmation and then find out that your transaction was actually on a fork and not on the main chain. This small chance of transaction failure is a

feature of Bitcoin and any other currency modeled on it, and is unappealing in a means of payment.

Cross-Exchange Price Discrepancy

One of the ways in which we determine value is through prices. However, in the case of cryptocurrencies it is difficult to obtain accurate prices. Most of the empirical work considers prices that derive from exchanges, although some trade occurs on-chain (that is, person-to-person). It is fair to say that the quality of data is not good. The exchanges are fundamentally unregulated, which leads to worries about price manipulation. And of course, the credit risk of the exchange itself is priced in the trades. It is important to remember that US equity markets are highly regulated. Thus, even though we hear of the occasional SEC enforcement action, these are the exception rather than the norm.

An early paper that documented discrepancy of prices across Bitcoin exchanges is Alexander Kroeger and Asani Sarkar (2017). They document persistent price differences across exchanges, which they attribute to credit risk. More recently, Igor Makarov and Antoinette Schoar (2018) estimate an astoundingly high size of arbitrage profits over a three-month period of \$1 billion US dollars. They present suggestive evidence that regulatory barriers such as currency controls may be behind these differences. This is also the view of young Jin Choi (2018), who present evidence for one country (South Korea) that currency controls are consistent with the perceived mispricing.

Illiquidity and Price Manipulation Evidence on market manipulation provides another source of concern when using the price to determine the value of cryptocurrencies. Reminiscent of the “pump and dump” type of schemes observed in penny stocks, there is evidence of price manipulation in Bitcoin. Specifically, Neil Gandal et al. (2018) use data from the Mt. Gox hack to observe that specific traders appear to have used the illiquidity of the market to induce a price run-up. They identify suspicious trades, and document that the exchange rate rose on average four to five percent per day during these periods. The data that they use is part of the interesting lore of Mt. Gox. Briefly, this exchange based in Japan was hacked, and the data (with an embedded Bitcoin mining virus) were posted in 2014 on the web. The data are detailed per Neil Gandal et al. (2018) and permit trade matching. Most importantly, the data include trader identifiers. The authors can therefore trace trading volume to individual trading strategies.

Similarly, John M Griffin and Amin Shams (2018) report suspicious trading activity in the purchase of Bitcoin using Tether on BitFinex. Briefly, Tether was designed as a digital asset that was not prone to price fluctuations. Tethers are in a class of what is known as “realcoins,” namely, coins that

are backed by “real” assets. In this case, Tethers are purported to be backed one-for-one with US dollars. The idea behind Tethers is clever. Real money flows through the banking system, which can be difficult for some customers to access for regulatory reasons. Tethers were purported to be equivalent to USD, but are accessible to all traders. After the leak of the Paradise papers, it became evident that Tethers were issued by the same individuals who operated the cryptocurrency exchange BitFinex. The USD that back Tethers are supposed to lie in a Taiwanese bank, but there do not appear to be regular public audits of the accounts. Against this backdrop, John M Griffin and Amin Shams (2018) look at the flows from Tether to several large cryptocurrency exchanges: BitFinex, Polloniex, and Bittrex. They find that Tethers are used to purchase Bitcoin after Bitcoin downturns, resulting in large increases in the price of Bitcoin.

Tao Li et al. (2018) document well-organized pump and dump schemes in cryptocurrencies. Briefly, a “pump and dump” scheme is a dubious stock market strategy typically associated with small stocks that have low liquidity. The purpose is to artificially inflate the price (typically through rumors) and then sell before the price collapses again. As the authors indicate, such schemes exist in cryptocurrency markets. Specific cryptocurrencies are targeted. A pump is announced on a Telegram chat to begin at a specific time. These machinations are successful: The authors document large price run-ups. In the first 70s after the pump begins, prices increase by 25% on average with a spike in volume.

Price manipulation is only possible if prices move in response to orders; in other words, if there is illiquidity in the market. First, we note that defining a “market” for cryptocurrencies is difficult as they can trade both on-chain (i.e., person-to-person) and on one of the many cryptocurrency exchanges.

The usual assets for which we have data are high regulated and so cryptocurrencies provide a unique experimental environment for researchers to investigate exactly how our regulatory regimes affect prices and outcomes. Some coins are issued by legitimate businesses, with Ripple being a prominent example. It is perhaps surprising that there is so little fraud in cryptocurrency markets.

If one is willing to accept that prices may be inflated and there may be episodes of illiquidity, a natural question for an individual investor is to determine if and how to invest in cryptocurrencies. The industry is ahead of this, in that institutional investors are starting to view cryptocurrencies as a natural asset class. There are now professional investors who manage portfolios of cryptocurrencies such as Polychain capital.

In addition, derivative securities are available on Bitcoin Futures traded on organized and regulated exchanges such as the CME and the CBOE. Hence, traders who do not want to or are not permitted to hold Bitcoin directly can still participate in speculation. The futures also make it possible for pessimistic traders to take short positions on Bitcoin.

Real-World Usage of Cryptocurrencies

A natural empirical question is how cryptocurrencies are actually used. Susan Athey et al. (2016) find that most buyers of Bitcoin are buy-and-hold investors rather than consumers who use Bitcoin for transactions. This makes it difficult to tie fundamentals (i.e., the value as a payment instrument) to cryptocurrency prices. However, we note that asset pricing for regulated securities (e.g., stocks) has not been particularly successful.

To get a sense for how “ordinary” investors use and think about Bitcoin, the Bank of Canada commissioned a usage survey. At the end of 2016, even though most Canadians knew what Bitcoin was, only 2.9% of the population actually owned any. These owners were predominantly younger males with a college education. Of those that did hold Bitcoin, the primary reason was general interest in new technology. Specifically, respondents did not find that it solved a pressing problem, but rather bought it out of interest. Scott Schuh and Oz Shy (2016) find similar results from an earlier survey of the US population.

Although users are not named on the blockchain, starting from addresses that had Bitcoin seized by the authorities for malfeasance, Sean Foley et al. (2018) use various statistical techniques to determine who uses Bitcoin and why. They estimate that about one-quarter of all users and just over 40% of transactions are affiliated with illegal activity. This amounts to about one-fifth of the dollar value of transactions and just over half of all the holdings. They also find that illegal trade moved out of Bitcoin as more opaque cryptocurrencies became available. Ironically, while the volatility of Bitcoin prevents the law-abiding from using it as a means of payment, the risk inherent in using cash means that those with criminal intent find it a useful way to receive payment.

3 Conclusion

While cryptocurrencies are currencies when used as payment instruments, they are also investible assets when investors use them to speculate and hedge. Cryptocurrencies can be used to raise funds for projects and goods, as an alternative to traditional fund-raising schemes such as IPOs. On the other

hand, cryptocurrencies face technological limitations, including the fact that maintaining the correct incentives in the blockchain is costly and inefficient. The market for cryptocurrencies is still underdeveloped and unregulated, and prices are subject to manipulation. Despite their promises, the usage of cryptocurrencies is limited and a large proportion is used for illegal activity.

References

- Abadi, Joseph, and Markus Brunnermeier. 2018. Blockchain economics. Mimeo Princeton University.
- Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia. 2016. Bitcoin pricing, adoption, and usage: Theory and evidence.
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta. 2018. The blockchain folk theorem.
- Bolt, Wilko, and Maarten R.C. van Oordt. 2016. On the value of virtual currencies.
- Breckenridge, Roeliff Morton. 1893. The paper currencies of New France. *Journal of Political Economy* 1 (3): 406–431.
- Budish, Eric. 2018. The economic limits of Bitcoin and the blockchain. National Bureau of Economic Research.
- Catalini, Christian, and Joshua S. Gans. 2018. Initial coin offerings and the value of crypto tokens. Manuscript.
- Chod, Jiri, and Evgeny Lyandres. 2018. A theory of ICOs: Diversification, agency and information asymmetry. Working Paper.
- Choi, Kyoung Jin, Alfred Lehar, and Ryan Stauffer. 2018. Bitcoin microstructure and the Kimchi premium.
- Cong, Lin William, Ye Li, and Neng Wang. 2018. Tokenomics: Dynamic adoption and valuation.
- Foley, Sean, Jonathan Karlsen, and Tālis J Putniņš. 2018. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?
- Gandal, Neil, J.T. Hamrick, Tyler Moore, and Tali Oberman. 2018. Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics* 95: 86–96.
- Garratt, Rodney, and Maarten van Oordt. 2019. Entrepreneurial incentives and the role of initial coin offerings. Working Paper.
- Griffin, John M., and Amin Shams. 2018. Is Bitcoin really un-tethered?
- Heaton, Herbert. 1928. The playing card currency of French Canada. *The American Economic Review*, 649–662.
- Howell, Sabrina T., Marina Niessner, and David Yermack. 2018. Initial coin offerings: Financing growth with cryptocurrency sales. Working Paper.
- Hu, Albert, Christine A. Parlour, and Uday Rajan. 2018. Cryptocurrencies: Stylized facts on a new investible instrument.
- Kocherlakota, Narayana R. 1998. Money is memory. *Journal of Economic theory*, 81(2): 232–251.

- Kroeger, Alexander, and Asani Sarkar. 2017. Law of one bitcoin price? Federal Reserve Bank of Philadelphia.
- Lee, Jeongmin, and Christine A. Parlour. 2018. Crowdfunding, initial coin offerings, and consumer surplus. Available at SSRN 3300297.
- Li, Jiasun, and William Mann. 2018. Initial coin offering and platform building. Working Paper.
- Li, Tao, Donghwa Shin, and Baolian Wang. 2018. Cryptocurrency pump-and-dump schemes. Available at SSRN.
- Liu, Yukun, and Aleh Tsyvinski. 2018. Risks and returns of cryptocurrency. National Bureau of Economic Research Working Paper 24877.
- Makarov, Igor, and Antoinette Schoar. 2018. Trading and arbitrage in cryptocurrency markets.
- Malinova, Katya, and Andreas Park. 2018. Tokenomics: When tokens beat equity. Available at SSRN.
- Nakamoto, Satoshi. 2008. Bitcoin: A peer-to-peer electronic cash system.
- Petukhina, Alla, Simon Trimborn, Wolfgang K. Härdle, and Hermann Elendner. 2018. Investing with cryptocurrencies: Evaluating the potential of portfolio allocation strategies. Available at SSRN 3274193.
- Raskin, Max, and David Yermack. 2016. Digital currencies, decentralized ledgers, and the future of central banking. National Bureau of Economic Research.
- Schuh, Scott, and Oz Shy. 2016. US consumers' adoption and use of Bitcoin and other virtual currencies.
- Sockin, Michael, and Wei Xiong. 2018. A model of cryptocurrencies. Manuscript.
- Zimmerman, Peter. 2018. Blockchain structure and cryptocurrency prices. Working Paper.



The Role of Utopia in the Workings of Local and Cryptocurrencies

Diane-Laure Arjaliès

AT THE VERY BEGINNING, THERE'S THIS DREAM.¹

1 Introduction

Although alternative currencies have long existed, the number of initiatives has exploded in the aftermath of the 2008 financial crisis. Alternative currencies, such as local and complementary currencies (LCCs) and cryptocurrencies, are currencies used as an alternative system of money to the one emitted and backed by nation-states, i.e., fiat currencies. Fiat currencies or “paper” currencies are government-issued legal tenders whose value is not dependent on any other asset (such as gold or other commodities). Fiat currencies include today’s Canadian dollar, the euro, or the British pound. In

¹ Source: <https://www.impact.eco/en/our-mission/>, accessed 18 December 2018. Full quote from the mission page of cryptocurrency impact Coin, “At the very beginning there’s this dream. Will you join us in making this dream come true?” Reproduced with permission from Impact Finance.

D.-L. Arjaliès (✉)
Ivey Business School, Western University, London, ON, Canada
e-mail: darjalies@ivey.ca

2018, it was estimated that there were more than 2700 cryptocurrencies² and 5000 LCCs worldwide (Savoie 2016).

Alternative currencies aim to transform global financial institutions by providing citizens with alternative monetary systems. Since the emergence of cryptocurrencies at the end of 2010s, their potential role in the capitalist system has been heavily scrutinized by governments and industries. Yet little is known about the goals and workings of such monies. This chapter aims to fill this gap by providing a framework that uncovers the ideals pursued by alternative currencies and the effects of those ideals on the production of money. Based on a comparative analysis of one LCC, *Sol Violette*, and three cryptocurrencies, *Bitcoin*, *Ǧ1 “June,”* and *impak Coin*, I show that despite targeting the same financial institutions, the utopia—or projects for society—pursued by alternative currencies can differ to a large extent. Uncovering these differences is essential if we want to understand the prevalence of this alternative movement and its potential consequences for our economies and societies. To do so, I analyze the valuation infrastructure of each currency—i.e., the mechanisms through which value(s) is attributed to the exchanges mediated by the currency (Friedland and Arjaliès, 2021; Helgesson and Muniesa 2013). Throughout, I demonstrate that the utopia pursued by the currency matters as much as its technology (e.g., blockchain) in shaping the production of money. This finding implies that to understand the workings of a currency, the focus of analysis should not be only on the technology used to issue new tokens—as it is often the case, but also on the criteria on which this issuance is based. Based on these results, I outline some implications for the social studies of financial technologies, their effects on our societies and their regulation.

I conclude that (some) alternative currencies do form a coalition of the unlikely, but not as the result of a strange bedfellows situation, in which individuals share common practices despite holding different views, but rather as the outcome of the opposite situation in which individuals share similar views but follow different practices. While all alternative currencies certainly aim to escape current monetary systems, they should not necessarily be interpreted as a manifestation of a post-capitalist society—where capital would be shared by all citizens and value created through knowledge (Drucker 1994), as it is often the case (Cohen 2017). Many alternative currencies actually derive their value from the ability of a handful of capitalists to control the means of production... of money. As such, most alternative currencies are actually a perfect instantiation of what capitalism is, i.e., a system where private entities

² <https://www.investing.com/crypto/currencies>, accessed 25 February 2020.

own the factors of production, including probably one of the most important factors in today's societies—capital itself.

2 What Is Money?

2.1 The Three Functions of Money

Economists tend to identify money³ through the roles it serves in society. Something can be considered as money if it acts as: a *media of exchange* with which to make payments; a *store of value* with which to transfer purchasing power (the ability to buy goods and services) from today to some future date; and a *unit of account* with which to measure the value of any particular item for sale. Different assets can play this role. For instance, Radford (1945) shows that cigarettes fulfilled these three functions for prisoners of war camps during the Second World War. Likewise, vouchers that represent a prepayment for goods and services from a specific supplier (or group of suppliers) can be considered as money.⁴ To function as a medium of exchange, there needs to be a system that enables transfers of value (i.e., a payment system). For any system other than the exchange of physical banknotes or coins, a means of recording the values stored (i.e., a ledger) is also required (McLeay et al. 2014).

2.2 What Is Blockchain Technology?

Non-cryptographic currencies rely on a centralized double-entry ledger (i.e., balance sheet) to record the value of money exchanged between individuals. Cryptocurrencies rely on a specific type of ledger, that is referred to as “blockchain.” Bitcoin is thus described as a digital ledger of monetary transactions in the form of a blockchain.⁵ A transaction is a message such as, “I, Paul, sent five coins to Ana.” To be validated, the message must respect three elements: (1) non-repudiation: Paul cannot later disavow it; (2) authentication: only Paul could have sent it, only Ana could have received it; and (3) immutability: the transaction cannot be altered later (Veneris

³ Note that a currency is often considered to be the unit of measure associated with money. For the sake of simplicity, I will use both terms interchangeably in the rest of the chapter.

⁴ However, their legal status generally differs from that of a banknote, since they do not usually legally entitle the holder to redeem the voucher.

⁵ The main innovation of Bitcoin was to combine the blockchain technology (which existed before 2008) with hashed digital signatures and timestamping in a peer-to-peer network.

2018). Blockchains aim to provide irrefutable proof that a set of transactions occurred between participants. They ensure that the three elements are respected using a peer-to-peer network that puts the transactions into a ledger certified in a decentralized manner. The ultimate purpose of the system is to avoid the problem of “double spending” when a given set of coins is spent in more than one transaction.

There are different types of blockchains. The blockchains used by business organizations are generally permissioned (between approved entities) and private. The one used by Bitcoin is an example of a permissionless (cryptographic approval) public blockchain. In the case of Bitcoin, transactions are grouped into blocks by miners, with these blocks forming hash-linked chains (i.e., the blockchain). Miners compete to randomly solve a cryptographic game to gain the right to create a new block. This puzzle is known as the “hash algorithm,”⁶ a process through which transactions are validated and new transactions are arranged into a sequence and time-stamped. A blockchain entry is hence both a receipt and a transaction. Instead of two independent ledgers (i.e., double-entry bookkeeping), the entry is recorded on a public, distributed, interlocked, permanent and transparent ledger. The transactions recorded in the ledger accumulate over time, which means that it is almost impossible to edit past transactions. Indeed, to deconstruct the blocks, the approval or “proof of work” would have to be entirely re-executed for every block added after the edited one, which is computationally extremely demanding (estimated at more than 51% of the computer power of the system—a total that amounts to 1% of the world’s energy consumption⁷). Bitcoin miners receive “block rewards” or “coinbase transactions” when solving the puzzle. Note that most cryptocurrencies now use “proof of stake” where the creator of the next block is chosen via various combinations of random selection and wealth or age (i.e., the stake).

⁶ The SHA (Secure Hash Algorithm) is one of a number of cryptographic hash functions. A cryptographic hash is like a signature for a text or a data file. The SHA-256 algorithm generates an almost-unique, fixed-size 256-bit (32-byte) hash. A hash is a one-way function—it cannot be decrypted back. This makes it suitable for password validation, challenge hash authentication, anti-tamper, digital signatures. Source: <https://www.xorbin.com/tools/sha256-hash-calculator>, accessed 25 February 2020.

⁷ Written Testimony of Arvind Narayanan, Associate Professor of Computer Science, Princeton University United States Senate, Committee on Energy and Natural Resources Hearing on Energy Efficiency of Blockchain and Similar Technologies, 21 August 2018. Source: https://www.energy.senate.gov/public/index.cfm/files/serve?File_id=8A1CECD1-157C-45D4-A1AB-B894E913737D, accessed 17 December 2018.

2.3 Money as Debt

Money is essential to the economy, but its nature has varied substantially over time. Despite the broad use of fiat currencies in contemporary capitalism, individuals still often believe that the value of the money they use is based on the reserves of the precious metal held by central banks—akin to commodity money. Due to this popular belief, today's mechanisms of creation of money are often misunderstood. The Bank of England (2014: 1) laments, "Money creation in practice differs from some popular misconceptions – banks do not act simply as intermediaries, lending out deposits that savers place with them, and nor do they 'multiply up' central bank money to create new loans and deposits."

The reality is that currencies issued by central banks only account for a very small amount of the money held by people and organizations (e.g., an estimated 3% for the United Kingdom [McLeay et al. 2014]). Most money in circulation (97%) actually consists of deposits with commercial banks, which are produced by banks themselves. These deposits are created when banks lend money to individuals, i.e., through debt. When a bank makes a loan, for instance by granting a mortgage to buy a house, the bank does not actually have the equivalent money in currency or even a proportion of the amount in its reserves. Instead, the bank credits the customer's bank account with a deposit the size of the mortgage. As such, the recording of the deposit is the very moment of production of money. Conversely, repaying bank loans destroys money, just as making loans creates it.

Central banks shape commercial banks' ability to lend money, either directly through the loan rates they charge the commercial banks or indirectly through the way their monetary policy affects economic activity. Such decisions influence the commercial banks' overall confidence in the ability of economic actors to reimburse loans, in other words, their ability to acquire economic value (e.g., through an increase in real estate values or through salaries) and the resulting capacity to create money. These decisions are reflected in the interest rate (plus any fees) charged by the bank, which dictates the amount that households and companies want to borrow, constraining in practice the amount of new money created. Thus, contrary to what individuals usually believe, the quantity of money that a commercial bank can create does not depend on the amount of reserves the commercial bank owns at the central bank, but on its confidence in the future economic well-being of its clients.

Commercial banks are the mechanism through which debts are centralized and cleared, meaning that government money is not more important

than bank money (Dodd 2016: 105). Currencies (banknotes and coins) are therefore better described as “debt” or “promises to pay” from the central bank to consumers and “bank deposit promises” from commercial banks to consumers. Bank deposits are in fact a record of how much a bank owes its customers. For the Bank of England, money is then “just a special form of IOU [I owe you], or in the language of economic accounts, a financial asset. [...] Financial assets are simply claims on someone else in the economy – an IOU to a person, company, bank or government” (McLeay et al. 2014: 1). What makes money special, according to the Bank of England, is that it is an IOU that everyone in the economy trusts and therefore accepts as a universal medium of exchange. As Dodd (2016: 134) summarizes, “Debt is no longer facilitating capitalism. It is driving it.”

2.4 Money as Social Relation

Considering money as an acknowledgment of debt has become the mainstream view in monetary theory (Dodd 2016: 93). Most sociologists of money complement this view by arguing that money is also a “social relation” (Ingham 1996, 2013), whose value depends less on it being underpinned by a commodity than on users trusting the system on which its issuance is based. Dodd (2016: 393) hence suggests shifting from a view of money as a “thing” with a “stable meaning and functions” to a “process that is inextricably social, inherently dynamic, and complex.” Zelizer (1989) was one of the first researchers to demonstrate that “a dollar is a dollar is not a dollar.” She showed that the money making up household income in the late nineteenth and early twentieth centuries was differentiated according to both gender and class differences. She named this process “relational earmarking” or the mechanism through which “people create, maintain, negotiate, or sometimes dissolve their social-economic relations by searching for appropriate matches among distinctive categories of social ties, economic transactions and media of exchange (Bandelj 2016; Zelizer 2012)” (Bandelj et al. 2017: 6). In other words, people associate each type of money with particular social and economic exchanges (for example, an individual expects to bring a bottle of wine rather than cash when invited to dinner) and, as I will show below, utopia. Money is thus transformed into an “object of desire” (Yuran 2014) not only for what it can bring individually but also collectively. With this in mind, sociologists of money interpret the creation of alternative currencies as a means of resisting a nation-state, whose sovereign privilege includes selecting the system of tokens it is willing to accept as payment of taxes (Cohen 2004; Knapp 1924).

The relational properties of money “enrich money from the inside through the social relations it makes possible” (Dodd 2017: 240). This enrichment unfolds through the interactions between individuals, where people investigate and make decisions based on their own fictional expectations (Beckert 2016: 14). The relational work involved in an alternative currency is intrinsically linked to the valuation infrastructure attached to the currency. The valuation infrastructure of the currency dictates how and what is deemed worthy of value, and comprises four dimensions (Friedland and Arjaliès, 2021): (1) institution: how the chosen value(s) are set up as ideals to pursue; (2) production: the mechanisms for producing such value(s); (3) territorialization: the selection of objects, practices or beings to which value(s) can be attached; and (4) evaluation: the evaluation schemes used to assess the production of (future) value(s). Value in this context is not limited to financial value and can include elements that individuals judge to be important in life, such as humanity, friendship, and pleasure. Like sociologists of money, sociologists of valuation suggest focusing on valuation as an action, rather than values as things (Hutter et al. 2015; Muniesa 2011). Valuation is “itself a relational, active process out of which something can hold as the sign (read ‘the value’) of something” (Muniesa 2011: 32).

Despite the essential aspect of valuation in the constitution of monies, the valuation process of money per se has attracted little interest. According to Ingham (1998), social theorists have tended to “sociologize” (p. 14) money with the result that most research has focused on the generation of “trust” and the social and political dimensions of monies at the expense of what money is intrinsically used for, i.e., to value. The author comments, “the approach should be balanced by a recovery of some responsibility for what are seen as ‘economic’ problems, such as inflation, the supply of credit, the determination of interest rates and so on” (Ingham 1998: 14). Central to the issues evoked by Ingham (1998: 14) is “the conception of the ‘real’ economy and its equally deficient anachronistic commodity theory of money.” This problem is of primary importance for monies like LCCs or cryptocurrencies such as *impak Coin* (see below) whose goal is to avoid “speculative value” (Davis 2018). Despite the increasing numbers of alternative currencies, the workings of such valuation processes are therefore relatively unknown. Understanding how these monies function could not only augment our knowledge of alternative currencies, but also contribute to the understanding of why the existing monetary system is repeatedly called into question.

3 What Are Alternative Currencies?

3.1 A Brief History of Alternative Currencies

During the Bretton Woods system (1944–1971), most of the allied nations' currencies (e.g., United States, Canada, Western Europe, Australia and Japan) were underpinned by gold, and were therefore known as “commodity money.”⁸ The United States actually used a gold standard for most of the late nineteenth and early twentieth centuries, prior to Bretton Woods. Unlike commodity money, the fiat money used today has no commodity-linked “intrinsic value.” It derives its value from currency users' trust in the government issuing the currency. Nonetheless, before the late nineteenth century and the expansion of the British monetary system model in the world (Knapp 1924), common national systems were the exception, with local currencies being the norm. Common currencies at that time were mainly used to compare and exchange local currencies. For instance, British colonies issued bills of credit with a predefined exchange rate with the pound sterling. These “colony currencies” existed from the end of the seventeenth century through to the late eighteenth century (Peacock 2014). For many centuries, European local lords could mint coins for exchange on their territories. Like the national commodity monies, the value of currency exchanges was most often dictated by the amount of precious metal (e.g., gold or silver) contained in each coin.

After the nineteenth century, however, most local currencies disappeared as a result of the rise of national and international monetary systems (Blanc 2006), only resurfacing from time to time, during economic crises or wars. For instance, during the Panic of 1907, the largest financial crisis in the US before the Federal Reserve was established in 1914, 145 US cities issued currency substitutes to the public to help their banks overcome the crisis (Yue 2015). What distinguishes today's western LCCs⁹ from their predecessors is the content of their utopia (Zelizer 2004). Previous LCCs were often used as a temporary solution for monetary systems whose currency value was undermined. Today's Western LCCs are mainly developed to fight the impersonal aspect of economic exchanges and inequalities embodied by fiat currencies. According to their proponents, LCCs could help build a better society by restoring rather than eroding human ties.

⁸ National currencies could be converted into US dollars, which were underpinned by gold reserves—the United States controlling two thirds of global gold reserves.

⁹ LCCs are used in some emerging countries due to the lack of a functioning national monetary system. This chapter only focuses on western LCCs projects.

Cryptocurrencies are a much more recent phenomenon. The first one, Bitcoin, was launched in 2009. A cryptocurrency is a digital currency transacted securely, transparently and peer-to-peer by means of cryptography (Hsieh 2018: 1). The launch of cryptocurrencies was motivated by a desire to eliminate the inefficiencies of the intermediated banking system that had prevailed so far in most capitalist societies. Such intermediated models rely on the ability of central banks¹⁰ and commercial banks to mediate the relationship between the supply and demand for money, either directly—through the issuance of coins and bills, or indirectly—through interest rates (see above).

The rest of the chapter focuses exclusively on contemporary Western LCCs and cryptocurrencies. These monies, however, are not the only types of alternative currencies existing today. “Community currencies,” with which LCCs are often associated, also comprise Local Exchange and Trading Schemes (LETS) that rely on barter and time-based money whose currency unit is time. LETS are usually not-for-profit organizations that aspire to alleviate poverty and social exclusion. The first LETS was created in Canada in 1983 to enable individuals to exchange goods and services in a situation where money was in short supply due to a local economic crisis. Time-based currencies involve valuing in time credits the hours spent on voluntary activities within a community (e.g., providing transportation or assisting an elderly person) (Meyer and Hudon 2017: 634). The Time Dollar currency (a time-based currency) was created in 1986 in marginalized areas of the United States to help rebuild social ties. Overall, community currencies aim to transform the nature of exchanges by promoting solidarity, cooperation and the development of sustainable development. They are therefore considered as local modes of resistance to the capitalist system and vehicles of social change.

3.2 The Utopia of (Alternative) Currencies

Since alternative currencies seek to offer imaginary alternative monetary systems for the societies in which they operate, they have also been described as “utopian monies” (Dodd 2016, 2017). Utopia refers to an imagined community or society that possesses nearly perfect qualities for its citizens. Utopia literally means a place “*topos*” that does not “*ou*” exist. Utopianism has often been associated with the abolition of money—a “dirty object” that organizations wanting to provide an alternative to capitalism usually seek to avoid (Farias 2017). Sociologists of money instead argue that utopianism is

¹⁰ The first central bank, the Bank of Amsterdam, was founded in 1609 in the Dutch Republic. A central bank’s mission is to promote the good of the people of its nation by maintaining monetary and financial stability.

actually implicit in all monies; each monetary system embodying an idealistic social project (Dodd 2016). Beckert (2016) further contends that any economic endeavor comprises “fictional expectations” that are essential to the workings of the capitalist system. Fictionality is the “inhabitation in the mind of an imagined future state of the world and the beliefs in causal mechanisms leading to this future state” (Beckert 2013: 219). Arjaliès and Durand (2019) thus show that each financial product is actually anchored in specific “morals” that dictate and shape market and human behaviors.

Utopia has always accompanied capitalism and its monetary system. After the French Revolution of 1848, Proudhon’s *Solution of the Social Problem* (Proudhon 1849) set out his project for two banks, a Bank of Exchange and a Bank of the People, which would give power to workers instead of financiers. Proudhon aimed to transform money into a medium of exchange rather than a means of amassing capital. He wrote, “Utopia needs for its realization capital accumulated, credit opened, circulation established and a prosperous state” (Proudhon 1849: 45 cited by Dodd (2017: 235)). The euro followed another utopia. Although the currency is now criticized for its lack of political and social vision, it was originally conceived as a way to help maintain peace and create cohesion in a region that had been bruised by war and occupation for years (Feldstein 1998; Kaelberer 2007). The euro envisioned a Europe that would stand together for better or for worse.

The aim of today’s LCCs is inherently “relational” (Zelizer 2012) in the sense that their ultimate purpose is to reconnect human beings to each other and to their territory. Money is not perceived as a thing, whose face value would derive from its intrinsic qualities (such as gold) but as a process, whose value comes from the mutual engagement of currency members. Such utopia is not the one shared by most cryptocurrencies whose members, in contrast, praise the non-specificity of money (Simmel 1904). Cryptocurrencies like Bitcoin indeed enable its members to not depend on anyone in particular, which means that no one has power over anyone. This non-specificity transforms money into a foundation for individual freedom.¹¹

¹¹ I am grateful to Jean-Philippe Vergne for raising the difference between non-specificity and anonymity. Cryptocurrencies’ members like Bitcoiners praise their freedom to exchange with whom they want (i.e., without the need to rely on existing financial institutions), but do not necessarily do it anonymously. There are in fact many social relations in most cryptocurrencies’ communities.

3.3 The Four Alternative Currencies Under Study

The following sections compare the utopia and valuation infrastructure (Friedland and Arjaliès, 2021) of four different types of alternative currencies—one local and complementary currency, *Sol Violette*, and three cryptocurrencies, *Bitcoin*, *ǧ1 “June”* and *impak Coin* (cf. Table 1 as overview).

Sol Violette (SV). SV was initiated in the city of Toulouse, in South West France, at the end of 2009 and officially launched in May 2011. This LCC was named the “Sol Violette,” after the violet, an emblem of the city. SV

Table 1 The valuation infrastructure of Sol Violette, Bitcoin, ǧ1 “June” and Impak Coin

Alternative currency	Institution: Values pursued	Production: Creation of values	Territorialization: Selection of objects, practices, beings to which values are attached	Evaluation: Schemes through which (future) produced values are assessed
Sol Violette	Encouraging local and humanized exchanges on a territory	<ul style="list-style-type: none"> – Pledging – Melting – Paper banknotes – Ethical merchants 	<ul style="list-style-type: none"> – Specific territory – Members sharing common values – Specific goods and services 	<ul style="list-style-type: none"> – Use of currency – Soft domestic product – Education of citizens – Ongoing resistance
Bitcoin	Disintermediating financial exchanges to re-empower individuals	<ul style="list-style-type: none"> – Mining using blockchain technology – Machine and social consensus 	<ul style="list-style-type: none"> – Not tied to a physical place – Anyone, any good and any service – Exclusive governance structure 	<ul style="list-style-type: none"> – Use of currency – Financial value—compared to other currencies
ǧ1 “June”	Valuing human beings for what they are, i.e., human beings, and not as productive forces of capitalism	<ul style="list-style-type: none"> – Creation of Universal Dividend using blockchain technology – Web of Trust 	<ul style="list-style-type: none"> – Ad hoc physical place – Living human beings who know each other – Any good, any service 	<ul style="list-style-type: none"> – Use of currency – Education of citizens – Ongoing resistance

(continued)

Table 1 (continued)

Alternative currency	Institution: Values pursued	Production: Creation of values	Territorialization: Selection of objects, practices, beings to which values are attached	Evaluation: Schemes through which (future) produced values are assessed
Impak Coin	Encouraging the development of an impact economy	<ul style="list-style-type: none"> – Creation of rewards using blockchain technology – Impak ecosystem 	<ul style="list-style-type: none"> – Members sharing common values – Specific goods and services – No “physical” place but a (virtual) marketplace 	<ul style="list-style-type: none"> – Use of currency – Transactions in the impact economy – Education of citizens

is a not-for-profit association that was founded with the monetary help of organizations from the social economy sector. The individual members (consumers) of the SV are called “Solistes” while the organizational members (producers and service providers) are named “Prestataires” (member organizations). In 2018, the LCC comprised 6 employees, 2400 members (Solistes) and 270 member organizations.¹² In 2015, a total of 172,546 SVs¹³ (equivalent in euros) had been exchanged and the “*Produit Intérieur Doux*” (PID), or Soft Domestic Product, generated was 287,130 euros. The permanent team is responsible for managing the community of member organizations and Solistes. Their activities include educating through events, diffusing the model by recruiting new members and building relationships with local institutions and other LCCs. The Sol Violette is considered by other European LCCs as a model for LCCs. The fact that it has been able to survive and maintain its democratic vision for over twelve years is considered evidence of its success.

Bitcoin. Bitcoin is probably the best-known cryptocurrency in the world. The Bitcoin white paper (Nakamoto 2008), which describes its functioning, was published in 2008 by Satoshi Nakamoto, a pseudonym for the lead (team of) developer(s). One of the key motivations for creating Bitcoin was to remove financial intermediaries, e.g., banks that charge fees to transfer

¹² <https://www.sol-violette.fr/indicateur/lister>, accessed 15 January 2019.

¹³ As a matter of comparison, this amount is what Bitcoin processes each second.

money between countries or a government-influenced central bank that uses taxpayer's money to bail out bankrupt private banks. In the aftermath of the 2008 financial crisis, the founder(s) of Bitcoin hoped to create a disintermediated peer-to-peer system based on blockchain technology that would put the financial system back into the hands of individuals, instead of financial institutions. In the nine years following its creation, Bitcoin's market capitalization increased from zero to US\$300 billion (maximum), with a total of 18,237,250 Bitcoins in circulation in 2020.¹⁴ Bitcoin has benefited from a high degree of (social) media interest. Its analysis by social scientists, however, has remained scarce, particularly from a sociology of money viewpoint, with some notable exceptions (see Bjerg 2016; Dodd 2018; Maurer et al. 2013; Swartz 2018; Vergne and Swain 2017; Vidan and Lehdonvirta 2019).

Ĝ1 "June." Ĝ1 (pronounced "June") is a "free" cryptocurrency launched in France in 2017. Free means that its issuance is neither based on any goods nor any production, but co-issued by all individuals alive in a community.¹⁵ This cryptocurrency aims to combine the human connection sought by the LCCs and the use of blockchain technology to produce money. The project is gaining momentum, but remains small in terms of members (4345)¹⁶ and volume exchanged.

Impak Coin. Impak Coin (MPK) is a combination of an LCC and a cryptocurrency. Launched in Canada in 2016 and currently in its beta version, it uses a private blockchain technology whose production of money is generated by transactions conducted in "impact" organizations that are members of the impak ecosystem. The goal of impak Coin is to recouple the value of money with the "impact economy," thereby contributing to the creation of a stable financial system at the service of a sustainable real economy. The company explains, "The impak Coin (MPK) merges the benefits and functionality of complementary currencies, rewards programs and new decentralized virtual currencies. The aim of the MPK is to encourage and reward the use of money into the impact economy—a 'real' economy composed of businesses that generates social and environmental positive impact." (source: interview)

Note that the four currencies differ to a large extent, in terms of types of organizations, scale, level of development and usage. I selected them since they embody different utopias while using, for some of them, the same type of technology. The sections are meant to be read in order, since the

¹⁴ <https://coinmarketcap.com/currencies/bitcoin/#charts>, accessed 25 February 2020.

¹⁵ <https://bravenewcoin.com/insights/how-blockchain-is-reshaping-our-economic-environmental-and-social-orders-ptIII>, accessed 10 January 2019.

¹⁶ <https://duniter.fr/g1/#/app/wot/map?c=46.6042:3.6475:6>, accessed 25 February 2020.

analysis of each currency builds on the previous one(s). With the exception of Bitcoin, whose founder(s) is/are anonymous, the founders and some users of all the currencies under study were interviewed for the purposes of this research (2017–2019). I also interviewed individuals who identified themselves as Bitcoiners, i.e., members of the Bitcoin community. These interviews enabled me to gather more information on the workings of each currency and better understand the motivations and ideals pursued by their founders and members. The analysis of Sol Violette also benefited from input obtained through an 18-month participative observation conducted by Delphine Gibassier as part of an ethnographic study.¹⁷ The descriptions below are based on my own interpretation of the interviews, observations and documentary evidence gathered on each currency. They should be ascribed neither to the interviewees nor to the organizations involved in each currency. Note that I have never possessed any of those currencies nor been involved in the organizations attached to the latter. The descriptions are schematic at best and focus exclusively on the utopia and valuation infrastructure of the monies.

4 Using Local and Complementary Currencies to Value Local and Humanized Exchanges: The Example of the “Sol Violette”

4.1 Institution of Value(s)

The Sol Violette (SV) was created “to encourage actors of the economy to produce and consume locally with respect for humans and nature, placing citizens at the heart of the project” (2011 SV Experimentation Report). SV members notably believe that their local currency could help address the “global crisis” that followed the 2008 financial crisis. In 2017, 84% of Solistes stated, “I adhere for ideological reasons, out of the conviction that finance is rotting society” (Internal Questionnaire, 2017). Their collective enemy is the “financialized economy” and its associated “speculation.” SV members want to put finance “at the service of the real economy” (2012 SV Report). They lament that most of today’s money is issued and managed in a private and competitive framework, through financial and banking institutions. According to SV members, the privatization of finance has led to our society being dominated by profit maximization. They explain that “Money

¹⁷ Data currently being analyzed with Delphine Gibassier and Tima Bansal.

is a public good, a necessity for all, but it is managed as any other private good, such as when individuals need to pay interest to obtain money” (2014 SV Report). In 2017, the SV disclosed its “wheel of values,” values that need to be pursued through the use of the currency. The wheel comprises eight values: relationships, life, proximity, democracy, ethics, citizenship, trust, and solidarity.

4.2 Production of Value(s)

The SV relies on a sophisticated and ambitious valuation infrastructure that comprises many different aspects. We describe the main mechanisms below.

Pledging. The principle of pledging (*nantissement*) is the first pillar of the SV’s operations. It involves going to one of the two partner banks to exchange euros for SVs, thereby taking money out of the traditional and speculative spheres of the economy and placing it in the “real economy.” Pledging money is described as an act of “*nantissement*” (“*nantis*” means “well-off” and “pledging a collateral as guarantee for a debt contract”) that “reorients money towards the productive sphere at the expense of the speculative circuits” (2011 SV Experimentation Report). SV members explain that “What is at stake is the re-localization of monetary exchanges, 98% of which are currently conducted on financial markets” (2013 SV Report) To incentivize members to pledge euros, the SV increases the value of each euro exchanged by 5%, meaning that 1 euro equals 1.05 SVs. The goal is to convert as many euros as possible into SVs.

Melting. The second pillar of the SV valuation infrastructure is melting (*la fonte*). Melting is a principle that causes a currency to lose value if it is not in circulation, e.g., a loss of 5% after three months of non-circulation. Melting is intended to increase the currency’s speed of circulation. According to SV members, “wealth is not created when money is accumulated but when it circulates” (January 2010, citizens’ meeting). By limiting hoarding, the SV hopes to multiply the speed of circulation of the SV by five, compared to the euro. According to SV members, this speed of circulation would create “five times more wealth while respecting humankind and nature.”

Paper banknotes. To encourage members to increase their use of SVs, the SV team debated switching to a digital version of SVs. SV members, however, believed that a “real economy” needs “real money” and therefore physical paper banknotes should not be abandoned. Paper banknotes are an “expression tool” that often triggers questions from non-members (2014 SV Report).

Selecting ethical merchants. To become an SV member organization, it is not enough to be a local economic actor. Organizations must either belong to the social and solidarity economic sector or prove their willingness to engage socially and environmentally for the benefit of the territory (assessed through a questionnaire). The SV's permanent team aims to guarantee that merchants' values are aligned with those of the currency and the political project it pursues (i.e., fighting the financialized economy). Each SV bank note states "Transform your money into a ballot!"

Empowering citizens. The LCC aims to transform citizens into "money experts" so they could introduce money into their daily conversations without being overwhelmed or questioned. To achieve such an outcome, the SV team aims to transform the LCC into a tool of "popular education." With this in mind, the team organizes open information meetings in different locations across the region with the hope to reach as many people as possible. These debates between (potential) members are essential for the LCC to fulfill its educational function. It is through these debates that individuals can reflect on and learn about the role of money in society. It is also a way for members to "feel responsible and own the project" (2011 SV Experimentation Report). The SV team also creates communication tools, such as pedagogical booklets that explain the benefits and workings of the SV, to transform what they judge to be a complex project into "something accessible." Money for SV members is therefore described as a political act.

In order to put money back on the political agenda, it is absolutely essential that citizens take ownership of it [money]. They need to be interested in how it works and in its many dysfunctions. It is at this moment that a local currency becomes a unifying educational tool. (SV 2013 Review)

4.3 Evaluation of Value(s)

Evaluation of SV stock (le stock). SV members use two main indicators to evaluate SV circulation. The first indicator involves assessing the "stock" of SVs in each member organization. The probability of exchanges taking place in SVs is directly linked to the diversity and abundance of SV supply and demand. The project's main difficulty is to achieve a level that ensures fluid exchanges.

Evaluation of SV leakage (la fuite). The second indicator used by SV members to assess the circulation of SVs is the absolute number of SVs converted back into euros. If member organizations or Solistes convert SVs back into euros, this means they have not been able to spend that amount in the SV circuit. SV members tend to avoid converting SVs because 5% of the

face value is retained by the SV (100 SVs become 95 euros) that reinvests it into some community projects.

Measuring the Soft Domestic Product (Produit Intérieur Doux - PID). The SV is expected to create wealth: employment, ethical values, and respect for humans and products. To assess the wealth creation generated by the LCC, SV members attempt to measure the PID “*Produit Intérieur Doux*” (Soft Domestic Product). The PID is defined as “the production sold that is respectful of humans and nature” (2011 SV Experimentation Report). The PID is meant to replace the PIB (*Produit Interieur Brut* or Gross Domestic Product). Transactions in SVs correspond to virtuous exchanges and contribute to the growth of the PID.

4.4 Territorialization of Value(s)

The SV is a local and complementary currency to the euro. It is local because the currency’s working principles promote local shops, short distribution channels and new forms of entrepreneurship orientated toward the sustainable development of the Toulouse region and its surrounding areas. The territory where the SV is used has increased continually since its launch. In 2012, SV members agreed that SVs could be used within a 100 km radius of Toulouse, provided that the suppliers furthest from the city maintained strong economic relations with Toulouse in order to guarantee the smooth movement of the currency (2012 SV Report). In 2011, to increase the circulation of SVs, the SV team began trying to extend the use of SVs to public services. The SV team worked with France’s central bank, the *Banque de France*, to attempt to allow payment of local services and taxes in SVs. However, until the 2014 French Law on the Social and Solidarity Economy, no public institution was ready to allow payment in LCCs. Another problem evoked by SV members is the lack of proximity merchants, cultural services and craft services in the network. The density of the distribution channel is judged to be too “weak” in many districts of Toulouse. SV members find it extremely demanding to physically travel to member organizations.

Despite their efforts, the SV team therefore finds it difficult to mobilize people within the network. Both Solistes and member organizations judge the daily use of SVs to be difficult and unnatural. The SV team explains that “it requires a lot of time and constant efforts to extend and mobilize new member organizations” (2012 SV Report). Individuals initially find it hard to understand how the LCC functions. When they finally grasp the workings of the project and decide to become members, they struggle with

the subtleties of the valuation infrastructure. Employees of member organizations find it particularly hard to use two currencies. Additionally, not all employees support the SV, which raises issues within member organizations. Sometimes they do not want to communicate on the project or to receive salaries paid in SVs. To address this issue, the SV offers member organizations specific accounting support when they first join the network and continuous education and training are needed to maintain the dynamism of the network. Yet proponents of the SV continue to praise the benefits of the LCC, which maintains some pressure on the financial system by acting as “a thorn in its side.”

5 Using a Cryptocurrency to Value Scarcity and Exclusivity: The Example of Bitcoin

5.1 Institution of Value(s)

Bitcoin is often described by popular media as a fashionable, risky, and speculative asset whose value is based on irrational beliefs. Yet the utopia that gave birth to the publication of the Bitcoin white paper in 2008, by anonymous Satoshi Nakamoto (2008), was quite different from the extreme form of capitalism that Bitcoin now embodies. Bitcoin was actually created in the aftermath of the 2008 financial crisis in order to remove the financial institutions from the monetary system to (re)empower citizens. It aimed to achieve this ideal by using blockchain technology and its peer-to-peer disintermediated system. In doing so, Bitcoin solved three key problems: (1) banks' central authority to validate transactions; (2) the fees incurred and the length of such a validation process and (3) the lack of anonymity of the actors implied in these exchanges.

Bitcoin actually allows users to transfer money, notably across borders, at a fraction of the cost and much faster than international wire transfers that use international settlement services such as SWIFT (e.g., Bitcoin fees estimated at less than 1% compared to 9% for Western Union (Wang and Vergne 2017: 2). The currency also provides users with pseudo-anonymity (only the user's digital wallet is traceable) and the three functions of money (see above) without any central authority. For all these reasons, Bitcoin has often been described as the product of anarchists and techno-utopians animated by their desire to suppress dominant social institutions by autonomous machines (Dodd 2018; Maurer et al. 2013; Swartz 2018; Vidan and Lehdonvirta 2019). The reality is probably much more complex. Today's Bitcoin users are

incredibly varied; some of them actually belong to the financial institutions that were first targeted by the cryptocurrency. Many individuals involved in the launch of Bitcoin from 2009 to 2011 for instance started other cryptocurrencies.¹⁸ While the utopia was quite clear when first created, today Bitcoin's ideals are certainly diverse.

5.2 Production of Value(s)

A key factor that needs to be understood in the functioning of cryptocurrencies is the mechanism through which new tokens are actually produced. What matters is not the blockchain per se, which is after all only a digital ledger, but the assumptions on which the blockchain is built. In the case of Bitcoin, new (fractions of) tokens are issued each time a miner successfully validates a transaction. These new tokens provide miners with financial incentives to conduct the “proof of work,” which is a fundamentally costly and random process. This “competitive bookkeeping” is called “mining” in reference to gold mining.¹⁹ The miners all compete to solve a cryptographic puzzle, and the winner is logically the one (e.g., individual or data center) with the most computing power. To get their transactions mined first, Bitcoin's users also include Bitcoin rewards in their transactions, with miners picking the transactions with the biggest incentives. The whole idea of Bitcoin is that as time unfolds, the cryptographic puzzle or “hash algorithm” will be more difficult to solve. It means that each mining operation will be costlier, which will eventually lead to fewer miners. The overall number of Bitcoins is indeed expected not to exceed 21 million. Akin to a gold-based currency, the value of Bitcoin is therefore based on Bitcoin users' belief that the number of tokens will never exceed this fixed amount, with the digital scarcity created through blockchain replacing the scarcity of precious metals (Dodd 2018: 37).

While the blockchain technology ensures a transparent,²⁰ peer-to-peer and disintermediated system, the governance of the blockchain ensures that the coordination is collective and the result of community decisions. To be validated, all proposed protocol changes in the code supporting the

¹⁸ By January 2019 and since the creation of the cryptocurrency, Bitcoin developers and network developers have “forked” the code, starting a new cryptocurrency, more than 70 times. The most successful fork to date is Bitcoin cash.

¹⁹ There is however an important difference with gold mining. When the price of gold goes up, miners can adjust the supply upward to make it go down and vice versa. Bitcoin has a fixed supply. No matter what the price of mining is, 12.5 new bitcoin will be issued in the next ten minutes.

²⁰ For instance, all miners can be identified using the blockchain's record of IP addresses. Some share their resources to form “mining pools.” Who is “behind” the IP addresses, however, remains unknown. See <https://www.blockchain.com/en/pools>.

blockchain must be supported by 51% of the computing power held by miners. This coordination mechanism is obtained through changes made in the Bitcoin Improvement Proposal (BIP),²¹ a design document providing information to the Bitcoin community, or describing a new feature for Bitcoin or its processes or environment. These choices rely on online and offline discussions. According to Hsieh and Vergne (2018), Bitcoin is actually one of the first “Decentralized Autonomous Organization” (DAO), i.e., “non-hierarchical organizations that perform and record routine tasks on a distributed, cryptographically secured, public ledger; and that rely on the voluntary contributions of their internal stakeholders to operate, manage and evolve the organization through a democratic consultation process” (Hsieh 2018: 1). According to the authors, DAOs coordinate tasks through a decentralized “machine consensus” (i.e., blockchain) that integrates a “social consensus” through a new class of paid stakeholders, namely, network validators (ibid.: 95). As we will see below, however, the democratic aspect of the system is questionable, as only the few individuals able to understand the complex workings of the currency can be involved in its governance.

5.3 Evaluation of Value(s)

In 2017, it was estimated that 450 developers contributed regularly to the code, that 200,000 transactions worth on average \$3,500 were processed every day, and that more than 11 million user accounts, known as “Bitcoin wallets,” existed (Hsieh 2018: 17). Bitcoin is therefore used as a system of payment. It is not yet used broadly as a unit of account, since many organizations still refuse to authorize the inclusion of cryptocurrencies in their balance sheet. Bitcoin is nevertheless by far the largest cryptocurrency in terms of market capitalization. Interestingly, however, this value is not given in Bitcoins, but US dollars. The ability to exchange Bitcoins for a fiat currency, and other cryptocurrencies, is actually essential to the valuation of the currency. The mechanisms of this valuation are unclear. Some researchers argue that the financial value of cryptocurrencies results from a mechanism of supply and demand, notably nurtured by (social) media effects, while others instead suggest that such value is intrinsic and depends on the technological innovation supporting the currency. According to Wang and Vergne (2017), for instance, cryptocurrencies do not behave like traditional currencies and do not follow the Quantity Theory of Money (Fisher 1911), according to which an increased supply should, *ceteris paribus*, lead to lower prices—and lower

²¹ <https://github.com/bitcoin/bips>, accessed 15 January 2019.

returns (Wang and Vergne 2017: 2). What is clear, however, is that the “trust” of Bitcoin users in the valuation infrastructure supporting the workings of the cryptocurrency, i.e., machine and social consensus, and to a certain extent in the fiat currencies to which their value is attached—hence the very financial system it aims to fight—are essential to its workings.

5.4 Territorialization of Value(s)

Bitcoin is not linked to any physical place. Bitcoin involves several stakeholders, the main ones being its miners, coders, and users. According to the Bitcoin Foundation, an American not-for-profit corporation that promotes and supports the currency, but which actually represents only itself, Bitcoin is a way for the entire of society to regain its “financial freedom.” They explain, “People everywhere are revolting against their escalating debt yoke, spiraling living costs and the unequal distribution of resources and wealth. Society as we know it is crumbling, and for good reason. Our financial system is broken and people are looking for another way.”²² Yet Bitcoin does not solve the very problem of accumulation of capital—which has often been identified as one of the key problems in today’s capitalism (Piketty 2013). Those who already have capital are likely to be richer than those who do not possess such wealth. And although Bitcoin does not function as debt per se—there is indeed no debtor and creditor, the value of the currency does rely on the anticipation of the overall increase of economic wealth—hence indefinite growth of capital. Bjerg (2016: 67) explains: “A Bitcoin does not represent a claim on any particular debtor but rather a claim upon the whole ‘society’ of Bitcoin users. Bitcoin is credit money without debt.”

The main argument put forward by Bitcoin miners is the ability of the currency to fight censorship by giving control to anyone who wants to be part of the community. As a matter of fact, thousands of individual miners, hundreds of committed developers, and thousands of occasional contributors participate actively in the governance of the currency. Yet questions remain regarding who is actually able to understand and engage in such tasks. The same questions apply to the users and buyers of such currencies. Who is able to grasp the mechanisms of valuation of those new types of assets? These questions are similar to those that have been raised by other financial products, such as high-frequency trading (MacKenzie et al. 2012). High-frequency products have been authorized by financial authorities, on the basis that they were considered as part of a “fair” market, since everybody could in theory

²² <https://bitcoinfoundation.org/>, accessed 15 January 2019.

access such technology (provided they have the capital).²³ With this in mind, is the fact that everybody could “in theory” join Bitcoin enough to argue that fairness is maintained? Such questions have led some Bitcoin miners to leave the community they describe as a new form of techno-capital elite and join new cryptocurrencies, such as Ĝ1 “June” (see next section).

6 Using a “Free” Cryptocurrency to Value Human Life: The Example of Ĝ1 “June”

6.1 Institution of Value(s)

The project of Ĝ1 (pronounced “June”) started in 2008 in the aftermath of the publication of an open access book shared by online communities and proposing a new relative money theory (Laborde 2019)—hence before the implementation of Bitcoin. The actual launch of the currency was done by several programmers in 2017—some of them having played with other cryptocurrencies before, and some of them being part of groups that launched LLCs.²⁴ So far, the cryptocurrency remains quite marginal, although it is arousing increasing interest, both in Europe and overseas. The currency grew from 59 members in 2017 to 4345 at the beginning of 2020.²⁵ The members of Ĝ1 have different ideological and socio-professional backgrounds, including libertarian, anarchist, extreme left, and apolitical, but all share the idea that the current financial and political institutions are broken. For the members of Ĝ1, the main problem with cryptocurrencies like Bitcoin is that the latter reproduces the current financialized capitalist system. In particular, Ĝ1 members lament that money supply envisioned by Bitcoin is not distributed equally among humans, instead absorbed by the owners of the most powerful computers, who are often the individuals who are already wealthy in fiat currencies. The main goal of Ĝ1 is to provide an alternative way of producing money: through life. Money in the Ĝ1 system is co-produced by its members, through the sole facts of their existence and

²³ See for instance, the mission of the Canadian Securities Administrators: “To give Canada a securities regulatory system that protects investors from unfair, improper or fraudulent practices and fosters fair, efficient and vibrant capital markets, by developing a national system of harmonized securities regulation, policy and practice.” <https://www.securities-administrators.ca/our-mission.aspx>, accessed 13 January 2019.

²⁴ See <http://www.monnaie-libre.creationmonetaire.info/monnaie-libre-n4-openudc/> and <http://www.creationmonetaire.info/2011/08/openudc-standard-monetaire-respectueux-des-derniers-entrants.html> for further information on the history, accessed 31 December 2018.

²⁵ <https://g1.duniter.fr/#/app/currency/stats/lg>, accessed 10 January 2019.

age. Ĝ1 provide all members with a daily number of tokens in a way that all generations are equally served in monetary creation share. The accumulated quantity of tokens—expressed in Universal Dividends—is expected to be the same for each individual over his/her life (estimated at 80 years). When a member dies, the money he/she accumulated slowly dissolves while the monetary mass increases.

The founders of Ĝ1 decided to incorporate the idea of a basic income into their cryptocurrency, through the Universal Dividend (UD). A basic income consists in providing all citizens with a sum of money that allows them to live decently, and this unconditionally of who they are and what they do. The idea behind the basic income is that revenues should not be based upon the ability of people to be productive forces, in (capitalist) economic terms, but should simply result from the fact that they are members of that society. The UD is calculated based on the assumptions offered by the relative money theory book (Laborde 2019).²⁶ The relative money theory involves restoring symmetry in terms of space and time to fight two injustices: firstly, the fact that only some individuals receive money (e.g., miners); and secondly, the fact that the first movers in a cryptocurrency will benefit from more value than those who join later. The relative money theory relies on four economic freedoms²⁷: (1) the freedom to choose our currency system, because money should not be imposed; (2) the freedom to access resources, because we should all have access to monetary resources (and consequently economic resources); (3) the freedom to estimate and produce value, because value is purely relative to each individual; and (4) the freedom to trade with money, because we should not be limited by the available money supply. Free or “freedom-enabling” cryptocurrencies are expected to support a free economic system.

6.2 Production of Value(s)

The calculation and production of the amount of money distributed to each member relies on a blockchain. Ĝ1 relies on a public blockchain known as Duniter, which does not involve mining. New tokens are automatically issued by the blockchain, on a daily basis and in the form of a UD. The number of tokens present in the UD evolves over time according to a formula from the relative theory of money. The formula is: $UD = G (M/N)$, where $G = \ln (LE/2)/(LE/2)$. UD = Universal Dividend, G = Growth of Money Supply

²⁶ https://wiki.p2pfoundation.net/Relative_Theory_of_Money, accessed 17 December 2018.

²⁷ https://wiki.p2pfoundation.net/Relative_Theory_of_Money, accessed 3 December 2018.

(estimated at 9.22% per year), M = Money Supply (total amount), N = Number of Members, LE = Life Expectancy (estimated at 80 years). Since the value of the UD depends on the number of members in the network, the currency is said to be co-produced by the members of the network. Everyone and every organization can agree to be paid, and pay, in $\check{G}1$, but only members of the network can produce the UD. The essential element for $\check{G}1$ members is that money is neither created through debt nor through the accumulation of capital through inflation. The first members to enroll in the network should not benefit from the increase in value of the currency (as it is the case for Bitcoin). It is also crucial that the value of the currency is not dependent on the value of another (fiat/crypto) currency, since the goal is to provide an entirely new and alternative system of money.

To achieve such independence, the blockchain algorithm (based on the above formula) is programmed so that all members accumulate the same quantity of UDs over their lifetime (i.e., 80 years). Note that 3740 UD is the value any member accounts tends to reach if there is no transfer and only the UD accumulating: the older UDs losing value over time as the total monetary mass increases.²⁸ The maximum number of members envisioned for each currency ($\check{G}1$ can be replicated) is estimated at 1,000,000 (5,000,000 technically) with the money from the deceased being constantly replaced by money from the living. Members have little interest in accumulating tokens since the value of goods and services should not be expressed in absolute numbers (i.e., $\check{G}1$) but in UDs—a value that is constantly adjusted according to the total money supply available in the network and the number of members. Provided all members stay in the network for the same length of time, there is no money asymmetry between members, whether in terms of space or time.²⁹

Unlike Bitcoin or the Sol Violette where converting a fiat currency into the cryptocurrency is the only way to enter the system (or by paying for goods and services in Bitcoins), $\check{G}1$ does not require the use of any other form of currency. To produce the daily UD, there is only one requirement, namely to be a member of the $\check{G}1$ network. To achieve this, $\check{G}1$ uses a “Web of Trust” (WoT), which ensures that every member is actually a living human being and that no one receives two or more UDs. The WoT used by $\check{G}1$ is inspired by, although it is operationally very different, the open source software Pretty Good Privacy, which is an encryption program that provides cryptographic

²⁸ This value differs from the actual number of UD created, which is 29,200 (one per day). Note however that the value of a UD created at a young age is worth almost nothing when the member is 80.

²⁹ Note that if two individuals of the same age enter at different moments, they are unlikely to accumulate the same amount of money over their lifetimes.

privacy and authentication for data communication.³⁰ The WoT is woven by members themselves. The 59 members who minted the genesis block asserted that they knew each other personally and each then proceeded to “certify” the other members around them. When a new member joins, he or she is given the right to certify future members.³¹ However, the process cannot go on forever as the web has a maximum diameter. Each member has a set number of certifications he or she can grant and cannot be “too far” from other members.³²

While there is some cryptographic competition to calculate the nodes of the blockchain, there is no incentive in winning the race (unlike bitcoin). Indeed, nodes are calculated voluntarily by members, who can be compensated in UD_s through a voluntary redistribution service—although few actually ask for such compensation. To keep the costs of the calculation as low as possible, the blockchain automatically adapts the difficulty of the cryptographic exercise to the computing power of the individuals wanting to calculate the node. Ĝ1 developers conceive the system to favor small computers like raspberry Pi,³³ maintaining the number of powerful computers to only a few so that the difficulty is considered properly. As a result, the currency is quite cheap to produce, both energy- and resource-wise.

6.3 Evaluation of Value(s)

The cryptocurrency is quite recent and its success is therefore fairly difficult to assess. The number of new members continues to increase on a regular basis, but the constraints imposed by the Web of Trust necessarily slow down its expansion. It indeed takes some time for each new potential member to be validated by five existing members. As for the Sol Violette, potential new members are recruited through word of mouth and public information sessions. At the beginning, most individuals have a hard time understanding how money is created, both in the case of Ĝ1 and in the existing monetary system (i.e., through debt). When people discover, however, that all they need to do to produce UD_s is to enroll online with a wallet, most of them are thrilled and ask to join the network. Proponents of Ĝ1 consider it much easier to implement than an LCC, notably because it relies on a simple phone

³⁰ https://en.wikipedia.org/wiki/Pretty_Good_Privacy, accessed 17 December 2018.

³¹ <https://duniter.org/en/introduction-a-la-toile-de-confiance/>, accessed 17 December 2018.

³² <https://duniter.org/en/duniter-why-how/>, accessed 17 December 2018.

³³ <https://www.raspberrypi.org/>, accessed 10 January 2019.

application³⁴ and there is no need to convert fiat currency in a specific bank. Ĝ1 is also said to be a much more egalitarian system than either the Sol Violette or Bitcoin since all members receive the same number of DUs over their lifetime.

In practice, the use of Ĝ1 remains quite limited (e.g., purchase of vegetables from a farmer, services between members, second-hand purchases through open source community forums) but the number of organizations (e.g., restaurants) willing to accept Ĝ1 is increasing. All documents are open source and the founders and proponents of Ĝ1 are actively involved online and in various communities to encourage the creation of other free cryptocurrencies elsewhere in the world. As for the Sol Violette, educating individuals on the workings of money is also an important measure of the currency's success. The ability to provide individuals with a concrete means with which to fight inequality is another. Unlike the Sol Violette whose members do not believe that LCCs can have a large impact via a scale effect, LCCs being by essence local and small in terms of volume, the Ĝ1 members envisage this possibility. Although they are aware of their project's utopia, they believe that for the first time in history, a real and alternative system of money creation is being offered.

6.4 Territorialization of Value(s)

Although there is no territory attached to the currency per se, the Web of Trust and the requirement to know five members of the network in order to produce the UD effectively bind the currency to a specific physical place. As with the Sol Violette, the Ĝ1 is expected to help build stronger ties between people by encouraging real transactions and discussions between members of the network. Yet there is no control as to what a “good” or “ethical” person is and consequently anyone can be part of the network, unlike the Sol Violette, which requests that their ethical merchants live within a specific territory. Likewise, any good and service can be bought or sold with Ĝ1.

The use of blockchain technology is also of secondary importance, it is described as a means to an end, a technological revolution that renders possible what was impossible before. In this respect, the founders of Ĝ1 strongly distance themselves from Bitcoin and other cryptocurrencies. They label themselves as part of the “free currency” movement, which does not preclude any specific form of technology. What matters to the proponents of

³⁴ A simple SMS service under development, as well as two physical paper systems too (paper wallets).

Ğ1 is the ability to use a currency to value human beings for what they are, i.e., human beings, not for what they could bring in capitalist terms. They explained, “As you’ll have gathered, Duniter’s main difference [compared with Bitcoin] is that **it puts humans first**.³⁵ The money creation is taken care of by humans themselves, not by machines or private corporations. Humans meet and validate each other’s capacity to create the money. In the best of worlds, each member would have a calculating node giving each and every one a say in the growth of the network and in future decisions. To avoid confusion between the protocol and the currency, we’ve distinguished the two, the former being Duniter and the latter the Ğ1. **Tech and economics are two very different subject matters!**”³⁶

7 Using a Cryptocurrency to Recouple Financial Valuation with the Real Economy: The Example of Impak Coin

7.1 Institution of Value(s)

Unlike previous alternative currencies that stemmed from grass-roots movements, impak Coin was created by a private corporation, impak Finance, with the explicit goals of making profits and having an impact. As with other currencies, the motivations of its founders were also triggered by the 2008 financial crisis and involved creating a better economic and financial system. The team that founded impak Coin is composed of seasoned and successful entrepreneurs who made a lot of money during the dot-com bubble. Although they come from the “technology” side, they have a good knowledge of the financial system, notably from the mergers and acquisitions and public listings in which they have been involved over their careers. This experience of raising capital nevertheless gave them the impression that a small elite was in charge of the entire financial system, distributing capital in a way that favored their own individual wealth at the expense of the productive real economy sustained by entrepreneurs.

In the aftermath of the 2008 financial crisis, one of the founders began to closely study the international monetary system and existing alternatives. During his research, he discovered that one bank—Triodos Bank—had lost

³⁵ Emphasized in the original.

³⁶ <https://duniter.org/en/duniter-why-how/>, accessed 17 October 2018.

no money during the crisis because all its assets were invested in “real economy” companies that pursued a triple bottom line approach (i.e., People, Planet and Profit). Once he had realized this fact, the founder started thinking about how he could use his technological knowledge to create a system enabling money to be spent, saved, and managed by organizations that support a real economy generating positive externalities. Convinced that this project could work, he and other seasoned entrepreneurs launched *impak Finance* in 2016, with the idea that this new venture had to be meaningful and a change maker in the world—their legacy for future generations.

In order to “test” their model, the team leveraged their network and organized a dozen “fun” design thinking workshops that gathered key representatives from the financial sector and tech companies. Participants were asked to imagine the bank of tomorrow. At the end of each workshop, participants realized that their “dream bank” was actually a 95% fit with the *impak Finance* project (source: interview with founder). Through this collective engagement, the team succeeded in raising CAD\$1.5 million in equity crowdfunding. In January 2017, *impak Finance* recruited its first employers, ready to become the first “neo-bank” whose cryptocurrency would be used to generate a social impact in the real economy. In September 2017, *impak Coin* closed their first Initial Coin Offering (ICO),³⁷ the first ICO legally authorized in North America, with a total of MPK 1,690,626 issued for CAD \$1,414,860 (89 countries, 2266 investors). With an additional CAD\$1.1 million secured through angel investors, *impak Finance* succeeded in raising CAD\$4 million over 18 months.

7.2 Production of Value(s)

The venture, which is B-Corporation certified,³⁸ is still very recent and most of its resources have been dedicated to fashioning the *impak* “ecosystem.” In fact, at the beginning of 2020, the MPK mobile application was still in a beta version. The reason for this delay is that the success of *impak Coin* relies on the ability to offer a large “market place” composed of “impact organizations” selling goods and services in line with the pursuit of a triple

³⁷ Investors invest money in return for a token, here the *impak Coin*, with the hope that the value of the token will increase in the future. More information on the ICO can be found at: <https://ico.bench.com/ico/impak-coin>, accessed 17 December 2018.

³⁸ Certified B Corporations are a new kind of business that balances purpose and profit. They are legally required to consider the impact of their decisions on their workers, customers, suppliers, community and the environment. This is a community of leaders, driving a global movement of people using business as a force for good. Source: <https://bcorporation.net/>, accessed 15 December 2018.

bottom line—all over the world. For this volume-based project to work, it is important that *impak* Coins could be used across a significant number of merchants, while also ensuring that these merchants deserve to be accredited “*impak* members.” The first two years of the venture hence involved creating a machine learning algorithm capable of assessing on a large scale whether the ventures eager to become members of the network could be qualified as “*impact*” ventures (i.e., as sources of positive externalities in the real economy), with qualified ventures then being enrolled. The criteria for measuring *impact* are based on the principles of the *impact management project*,³⁹ an international initiative that supports the sustainable development goals (SDGs).⁴⁰

The first strength of *impak Finance* is therefore their ability to provide a standardized verified “*impak profile*” that enables each citizen, investor or consumer to assess in a few minutes whether the organization from he/she wants to buy/invest contributes to an *impak* economy, notably through the pursuit of the SDGs. Once organizations become “*impak accredited*,” they can enter the *impak* ecosystem—made visible on a map in the application. Each time an *impak* registered user pays a member organization in fiat currency (or in *impak* Coins), the mobile application traces the “*impak*” transaction, and he/she receives MPK rewards (2–5% of cashback). *Impact* businesses are incentivized to join the network to (1) increase their sales within an extremely mobilized community of potential clients (both B2C and B2B); (2) have access to the free *impak* Coin “incentive” program; and (3) be visible from the *impact* institutional investors seeking *impact* qualified opportunities. The *impak* ecosystem also incentives qualified businesses to convince their own providers to join, as they will receive MPK cash back for each business-to-business transaction within the network.

While the Bitcoin blockchain issues new tokens when a block is mined, with *impak*, new tokens are created in the form of “MPK rewards” when transactions occur in the *impak* ecosystem. Fiat currencies can also be converted into MPK (as for the *Sol Violette*), but only to the extent that the governance structure of *impak Finance* authorizes it (see below). The blockchain technology employed is a privately commissioned, permissioned blockchain based on the Ethereum technology (i.e., smart contracts) developed under an open source licence by JP Morgan for its *Quorum*

³⁹ <https://impactmanagementproject.com/>, accessed 17 December 2018.

⁴⁰ The Sustainable Development Goals are the blueprint for achieving a better and more sustainable future for all. They address the global challenges we face, including those related to poverty, inequality, climate, environmental degradation, prosperity, peace and justice. The Goals interconnect and in order to leave no one behind, it is important to achieve each Goal and target by 2030. Source: <https://www.un.org/sustainabledevelopment/sustainable-development-goals/>, accessed 16 January 2019.

blockchain.⁴¹ In other words, *impak Finance* maintains control of the blockchain, there are no miners involved as for Bitcoin and no nodes calculated by its members as for *Ĝ1*. In due course, *impak Finance* nevertheless expects some selected partners to participate in the creation of nodes. The public keys, which are online “traces” of the transactions validated through the blockchain, will remain publicly available in order to enable future audits of all transactions conducted in the ecosystem.

Akin to the perfect money theory⁴² developed by Simmel (1904), the goal of *impak Finance* is to base the supply of money on the quantity of goods and services exchanged in the real impact economy. When an organization is paid in *impak Coins*, it can use them in the ecosystem to pay its suppliers (which increases its impact score) or convert them back into fiat currency for a fee of 2.5%. *Impak Coins* could also be used for peer-to-peer payments, with no transaction fees. Note, however, that MPK rewards cannot be converted into fiat currencies in order to increase the circulation of money (as for the *Sol Violette*). *Impak Finance* ultimately hopes to provide organizations with loans paid in *impak Coins* through microlending, crowdlending and peer-to-peer lending. This system would allow savers to earn interest in *impak Coins* while investing in the impact economy. The entire financial system would thus be re-routed and re-invented to support the financing and purchasing of goods and services with positive externalities, while reducing the number of financial intermediaries.

7.3 Evaluation of Value(s)

The intrinsic value of the *impak Coin* should derive from the growth (or decrease) of the worldwide impact economy and the social and environmental positive externalities of the impact economy it traces. To avoid any speculative behavior and guarantee that the value of *impak Coin* reflects the value of the *impak economy* on which it is based, *impak Finance* controls the conversion and value of the *impak Coin* in terms of fiat currency. The *impak Coin* is currently worth CAD\$1 and will hold the same value until the launch of *impak Coin* transactions in the *impak marketplace*. Akin to the central banks’ approach with commodity money and to ensure that the value of *impak Coin* remains “stable,” *impak Finance* have created an independent governance structure that will adjust the variables (e.g., % of money cashback) of

⁴¹ <https://www.jpmorgan.com/global/Quorum>, accessed 4 January 2019.

⁴² According to the perfect money theory developed by Simmel (1904), the stability of commodity money could be obtained through the maintenance of a just proportion—expressed in prices—between the total quantity of money in circulation and the total quantity of commodities on sale.

the mathematical function fixing the rate of conversion in terms of CAD\$ once a week, and buy or supply MPK with their own reserves (CAD, EUR and MPK initially) to adjust the quantity of money available in the system. Unlike Bitcoin and other “non-stable” cryptocurrencies, *impak* Coins will not be exchanged on other stock exchanges in order to guarantee that the value of the cryptocurrency is not dictated by other currencies or arbitrage opportunities. As for other cryptocurrency exchange places (and fiat currencies issued by central banks), *impak* Finance will also maintain ownership of all coins by owning all the “private keys” used by members to execute their orders on the blockchain (note that each order needs a public and a private key to be validated).⁴³

Impak finance is expected to make money on several dimensions—on fees (transactions, lending) but mainly through its system of impact assessment, scoring and traceability that will be sold to impact institutional investors and organizations eager to have access to qualified impact businesses and to evaluate their portfolio or supply chain’s impact. The cryptocurrency is described as follows, “Coded to support the impact economy development. [...] the first stable cryptocurrency designed to support the growth of the impact economy. By design, *impak* Coin will build loyalty, reward collaboration and encourage its holders to buy from *impak* accredited members.”⁴⁴ According to its founders, the success of *impak* Coin will be indicated by the quantity of money actually in circulation in the *impak* ecosystem. In December 2018, the *impak* Ecosystem said to comprise a community of nearly 9,500 citizens and tens of thousands of businesses and organizations. As with the *Sol Violette* and *Ĝ1*, *impak* Finance also seeks to educate citizens on the current monetary system and its (fatal) flaws through their use of the cryptocurrency.

7.4 Territorialization of Value(s)

The success of *impak* Coin will largely depend on its ability to create an ecosystem where individuals (consumers, savers) and organizations share the same desire to contribute to an impact economy. Like the *Sol Violette*, the goal of *impak* Coin is to use money as a way to create a real and sustainable

⁴³ Many cryptocurrency stock exchanges actually “own” the currencies present in their users’ wallets by owning the private keys (i.e. passwords) associated with the wallets. Many open source community members judge this ownership abusive and prefer to keep a “physical wallet” for their money (akin to a USB key). To be able to spend this money, however, the private key is needed, which means that even if someone finds the physical wallet, he/she will not be able to spend the money if she/he does not have the private key (or password). Likewise, if an owner loses the physical wallet, he/she actually loses the money kept in the wallet.

⁴⁴ <https://www.impact.eco/en/impak-coin/>, accessed 17 December 2018.

economy that emphasizes human connection and the preservation of nature. Unlike an LCC, however, this circuit is not linked to a specific territory and many of the merchants are actually online businesses. Like Ğ1, *impak Coin* also aims to provide an alternative to the debt model. Indeed, the MPK money supply does not depend on a potential future economic value, but on the transactions that occur in the real economy. However, the model does not address issues raised by the accumulation of capital (Piketty 2013), in the sense that the individuals who first acquired the tokens are very likely to benefit from the increase in value of the currency. Nevertheless, and unlike Bitcoin, the value of *impak Coins* does not derive (exclusively) from their exchange value (against other currencies) and their ability to function as a store of value (e.g., by building exclusivity through limiting the number of tokens). The transactions conducted in the real impact economy are indeed expected to anchor the financial value of the currency. The website explains, “An entrepreneur fuelled by a passion for technology, a sustainable development activist and precious, expert allies from complementary backgrounds. Add to that our growing, general sense of urgency in the face of immense global issues, and some major challenges in between, and we end up with an inspired team that took the bull by the horns. We decided to do our part. A little like all entrepreneurs who combine dreams and responsibilities. [...] At the very beginning, there’s this dream. Will you join us in making this dream come true?”⁴⁵

8 Discussion

Since the 2008 financial crisis, the number of alternative currencies all over the world has exploded. Yet, little is known about their goals and workings. This chapter aimed to fill this gap by providing a comparative analysis of the valuation infrastructure of one local and complementary currency (LCC), *Sol Violette*, and three cryptocurrencies, *Bitcoin*, Ğ1 “June” and *impak Coin* (cf. Table 1). The valuation infrastructure (Friedland and Arjaliès, 2021) of each currency describes how the value(s) pursued by each currency is instituted, produced, evaluated and territorialized. Based on this analysis, I could show that (1) despite targeting the same financial institutions, the utopia (or societal project) pursued by LCC and cryptocurrencies can vary to a large extent and (2) this utopia shapes the workings of such currencies as much as the technology used to produce money (e.g., blockchain). This framework

⁴⁵ <https://www.impak.eco/en/our-mission/>, accessed 18 December 2018.

could be applied to other currencies and innovations, thus contributing to a more systematic analysis of financial technologies. Through this analysis, I also illuminated the three functions of money (i.e., media of exchange, store of value and unit of account) performed by the four currencies, as well as their status as debt and social relation (cf. Table 2). As Table 2 shows, such functions vary and none of the alternative currency existing currently can be considered as equivalent to a fiat currency, since no country allows their taxes to be paid with such money. The analysis also indicates that the usage of money as a store of value, i.e., with an idea that its value will increase in the future, is associated with its role as debt. Lastly, it confirms that all monies are social relations. I will now elaborate on the implications of these findings for the studies of (financial) technologies, their impacts on society and their regulation.

8.1 On the Relationships Between Utopia and Technology

Bitcoin has been described as a “techno-utopia” (Dodd 2018: 42) whose goal is to use technological means to provide an alternative monetary system to the outdated banking system, described as full of political hubris, financial ineptitude and underlying criminality. As shown above, however, this focus on the technology is somehow misleading. In cryptocurrencies, the “social consensus” is as much important as the “machine consensus” (Hsieh 2018; Hsieh and Vergne 2018). Such currencies are therefore relational projects as much as LCCs are (Zelizer 2012). As Bjerg (2016: 62) explains, “Bitcoin has value as money only in so far as there is a community of users willing to accept Bitcoin in exchange for commodities, services, or other forms of money. Bitcoin does not have any intrinsic value.”

While not core to the sociology of money, the relationships between technology and utopia have attracted a lot of interest in other fields. Socio-materiality scholars have studied how material devices interact with social practices (Jarzabkowski and Kaplan 2015; Leonardi 2013; Orlikowski 2007). Social studies of finance and accounting scholars have shown that calculative devices are not neutral but instead shaped by the assumptions of their designers (Busco and Quattrone 2018; MacKenzie 2011; Miller and Power 2013). Strategists have also explored the relationships between the value proposition and the technological innovation of new business models and strategies—notably to know which one precedes the other (Aversa et al. 2015). The findings of this chapter are aligned with this body of research. Blockchain technology, as other technologies, is not “neutral” but instead

Table 2 The three functions of money of Sol Violette, Bitcoin, Ġ1 “June” and Impak Coin

Alternative currency	Media of exchange	Store of value	Unit of account	Money as debt	Money as social relation
Sol Violette	Yes, but within the SV community, although some conversion into fiat currency is possible	No, and to be avoided, as money should circulate. The value of the currency decreases over time if it is not being used	Yes, but within the SV community only, with some possibility to identify the currency as “voucher” into sovereign accounts	Yes & No, the value of money does not depend on the future availabilities of members to reimburse, although it is indirectly linked to their ability to do so in fiat currencies	Yes, money is used to connect like-minded people
Bitcoin	Yes, but within the Bitcoin community, although some conversion into fiat currency and other cryptocurrencies is possible	Yes, based on the foreseen digital scarcity of the currency, valued by stock exchange markets	Yes, but within the bitcoin community only, with some possibility to identify the currency as “intangible asset” into sovereign accounts	Yes & No, the value of the currency depends on the whole capacity of Bitcoin users to reimburse, not a particular debtor	Yes, the value of money depends on the belief of the community in it
Ġ1 “June”	Yes, but within the Ġ1 community only	No, and to be avoided, money should not support the accumulation of capital. The value of the currency stays the same over time (relative value)	Yes, but within the Ġ1 community only	No, the value of money does not depend on the future availabilities of members to reimburse	Yes, money is used to celebrate human life and relationships

Alternative currency	Media of exchange	Store of value	Unit of account	Money as debt	Money as social relation
Impak Coin	Yes, but within the impak community, although some conversion into fiat currency is possible	Yes, based on the expected increase of value(s) being produced within the impak community, valued by the community	Yes, but within the impak community only, with some possibility to identify the currency as "loyalty rewards" into sovereign accounts	Yes & No, the value of the currency depends on the whole capacity of impak users to reimburse, not a particular debtor	Yes, money is used to connect like-minded people

fashioned by the ideals of their creators and users. The technology shapes these ideals in return, as the currency cannot exist without the device that embodies it.

By uncovering the relationship between technology and utopia, the chapter demonstrates that some cryptocurrencies are more similar to some LCCs than to other cryptocurrencies. A cryptocurrency like *impak Coin*, for instance, is much closer in its functioning and utopia to an LCC like the *Sol Violette* than it is to a cryptocurrency like *Bitcoin*. Since alternative currencies all aim to challenge existing financial institutions, previous research has tended to group all cryptocurrencies and LCCs under one homogeneous group (Cohen 2017). This chapter shows that such analysis is both theoretically and empirically questionable. There is a huge diversity of utopias among alternative currencies—particularly cryptocurrencies—and such utopia is not dictated by the technology. Also, cryptocurrencies do not necessarily oppose LCCs. There could indeed be similar aspirations in both types of currencies. In other words, an LCC can theoretically use a blockchain technology, hence becoming a cryptocurrency as well, provided its use remains within a local context and in a way that does not search for the “non-specificity” of money (Simmel 1904) (i.e., the lack of mutual dependence between users). Uncovering these differences is essential to be able to understand the dynamics that inform the rise and development of alternative currencies in the world. If all alternative currencies aim to transform global financial institutions, the mechanisms through which they search to do it and therefore their potential impact vary to a large extent.

8.2 On the Ever-Changing Utopias

The rise of alternative currencies in the aftermath of the 2008 financial crisis has been fuelled by a strong rejection of the financial system across the entire political spectrum, as exemplified above. In this sense, LCCs and cryptocurrencies do form a coalition of the unlikely whose common basis is the search for another, more meaningful, monetary and economic system for our world. As explained above, however, what this utopia comprises varies considerably. Members of the currencies themselves can have different interpretations of the same utopia. For instance, some members of *G1* would love to use the cryptocurrency as a substantial source of revenues for the poorest members, while other members are mainly interested in the fair distribution of money between generations. Likewise, the utopia can evolve over time. *Bitcoin* is a good example. Not only individuals interpreted the utopia

described in the Bitcoin paper written by Satoshi Nakamoto (2008) differently, their interpretation evolved as the currency developed. In addition, and although white papers are broadly used, not all alternative currencies explain in explicit terms their utopias. On this dimension, *impak Coin* which is entirely managed by a private organization with a clear control of the currency clearly differs from Bitcoin which originated from a grass-root movement. As Frenkel (1977: 12) explains: “Money is not a consciously created artifact, but grows out of, reflects, and in turn affects the ever-changing relationships between individuals and the society which they compose.”

With this in mind, several questions can be explored by further research. Are some utopias more likely to support the development of the currency? Is the utopia likely to fade as the currency grows? Are alternative currencies having a loosely defined utopia more likely to evolve and adapt? Is there a systematic relationship between the type of technology (e.g., type of blockchain technologies) in use and the utopias pursued? Such questions echo some research conducted in the literature on social movements and institutional theory notably (Battilana et al. 2009; Benford and Snow 2000; De Bakker et al. 2013; McAdam and Scott 2005). However, it remains to know if alternative currencies can be theorized as social movements per se, notably due to the diversity of goals pursued and the variety of cryptocurrencies. Other scholars have thus preferred to describe cryptocurrencies like Bitcoin as (decentralized autonomous) organizations (Hsieh and Vergne 2018). As a matter of fact, most LCCs are actually non-for-profit organizations and *impak Coin* is a B-Corporation. Likewise, it remains uncertain whether alternative currencies succeed in transforming the institutions they target, to wit: global financial institutions.

8.3 On the Regulation of Alternative Currencies

The findings of the chapter have two main implications for the regulatory bodies in charge of alternative currencies, particularly crypto ones. Firstly, the chapter demonstrates that the utopia matters as much as the technology used to produce money. This observation is very important as current regulation relies quasi exclusively on the technology in use, at the expense of the type of value(s) pursued and produced. This analysis applies to the broad spectrum of financial technologies that have appeared over the past years (e.g., robot advising, artificial intelligence, crypto-assets, etc.). This chapter shows that none of these technologies, whatever complex and sibylline they appear, are “value-neutral.” Blockchain technology is multiple and the assumptions under which tokens are issued have considerable influence on the form and

content of the cryptocurrency. Uncovering a currency's utopia is therefore of primordial importance when attempting to understand its meaning and workings, thereby confirming the importance of envisioning money as "social relation" (Ingham 1996, 2013). Researchers and practitioners alike should therefore investigate the ideals as much as the technology when analyzing such innovations. Policy-makers, in particular, should consider the purposes of such innovations when making decisions regarding the added-value of these new activities for the functioning of markets.

Secondly, the chapter shows the importance of reflecting on the creation and distribution of value(s) supported by alternative currencies. Cryptocurrencies certainly appeared as a response to the increase of private accumulation of value by financial intermediaries, and this at the expense of the ultimate savers (Arjaliès et al. 2017). The decrease of transaction fees obtained through blockchain technology however does not necessarily mean that the distribution of value along the investment chain is fairer. Yet most of the decisions made by regulators seem to take the ability to lower transaction costs as the main criteria for deciding whether a financial technology is beneficial for the markets. For decades and in the aftermath of the Chicago School of Economics, the efficiency of markets was judged on the ability of a financial innovation to optimize supply and demand, often referred to as increasing the liquidity of markets (Fox 2009). Such rules dictated the usage of algorithms to determine the stock value (Muniesa 2007) or the legality of practices such as high-frequency trading (MacKenzie et al. 2012). As financial markets are shifting to passive investment, with the use of trackers that replicate indexes, investors increasingly make money on the ability to reduce transaction fees rather than finding opportunities in the markets. Throughout and with the rise of financial technology such as robot advising and blockchain, the regulatory bodies also seem to move from efficiency of markets to transaction cost theory to decide whether a new financial technology is appropriate (Fabozzi et al. 2014; Williamson 1991). This chapter shows that such approach does not help understand which value(s) is being produced, nor for whom. No regulation of this new phase of technological finance will be achieved if the discussion around the values that financial markets should pursue is kept being silenced. Technology shapes practice, as much as utopia shapes technology. And markets are not exempt from this relation.

9 Conclusion

To conclude and although the chapter does not cover this specific aspect, I want to acknowledge that many cryptocurrencies have been launched for financial reasons only. Those financial goals do not mean that there is no utopia, but that this utopia is likely linked to the ideals attached to a market-based economy (Arjaliès and Durand 2019). Hence cryptocurrencies do not only involve the techno-utopian anarchists that previous research has tended to depict but also include capitalists likely to use their wealth to appropriate the means of production of capital itself. While cryptocurrencies such as Bitcoin do prevent financial intermediaries from taking fees and power over transactions, it is unclear whether this re-appropriation of value is redistributed equally among those involved in the investment chain (Arjaliès et al. 2017). It is indeed very possible that cryptocurrencies are nurturing a new form of techno-financial elite that combines the source of capital and technical knowledge—an elite who might eventually be the same as the one in charge of current financial systems (since the latter are the ones who can buy computers and invest in crypto-assets).

Cryptocurrencies are therefore not (only) an expression of a post-capitalist society praising the sharing of capital and knowledge across all citizens (Drucker 1994), as previous research implied (Cohen 2017). While such observation applies to most LCCs, most cryptocurrencies are indeed probably better described as a form of capitalism where ownership of capital is obtained outside dominant social and economic institutions and their control, and this, thanks to new technological media. Yet this chapter also demonstrates that the same technology could be used for different purposes. The examples of Ğ1 and impak Coin show that citizens and entrepreneurs alike are willing to reinvent the society in which they live. Obviously, both projects are in their infancy and their success is as mythical as their dreams. As for LCCs, their members scramble for power and impact, but none of them has ever been lacking aspiration. In an economy of desire, what money wants remains utopia.

Acknowledgements I would like to thank all the founders and users of the alternative currencies who kindly opened the doors of their practices and utopia to my research. I am also very grateful to Delphine Gibassier who allowed me to use data on the local and complementary currency Sol Violette that she collected during an 18-month ethnographic study. I also thank my students, Michael King and Jean-Philippe Vergne for encouraging me to explore this fascinating (digital) world and sharing with me their knowledge on the topic. I also want to thank Paolo Aversa, Tima Bansal, and the participants of the 2019 Egos track “Dreams,

Fictions and Calculations: Imagined Futures in Organizational Life” organized by Klaus Weber, Jens Beckert and Brooke Harrington, for their precious feedback on previous versions of the chapter. Last, I want to thank Raghavendra (Raghu) Rau and Robert Wardrop, co-Directors of the Cambridge Centre for Alternative Finance, Judge Business School for supporting this project and making it come true. The interpretation of findings, however, is my own and should not be ascribed to others. Note that this chapter is part of a larger project on the rise of financial technologies and their impacts on society.

References

- Arjaliès, D.-L., and R. Durand. 2019. Product categories as judgment devices: The moral awakening of the investment industry. *Organization Science* 30 (5): 885–911.
- Arjaliès, D.-L., P. Grant, I. Hardie, D. MacKenzie, and E. Svetlova. 2017. *Chains of finance: How investment management is shaped*. Oxford: Oxford University Press.
- Aversa, P., S. Furnari, and S. Haefliger. 2015. Business model configurations and performance: A qualitative comparative analysis in Formula One racing, 2005–2013. *Industrial and Corporate Change* 24 (3): 655–676.
- Bandelj, N. 2016. Thinking about social relations in economy as relational work. In *Re-imagining economic sociology*, ed. P. Aspers and N. Dodd, 227–251. New York: Oxford University Press.
- Bandelj, N., F.F. Wherry, and V.A. Zelizer. 2017. Introduction—Advancing money talks. In *Money talks: Explaining how many really works*, ed. N. Bandelj, F.F. Wherry, and V.A. Zelizer. Princeton: Princeton University Press.
- Battilana, J., B. Leca, and E. Boxenbaum. 2009. How actors change institutions: Towards a theory of institutional entrepreneurship. *Academy of Management Annals* 3 (1): 65–107.
- Beckert, J. 2013. Imagined futures: Fictional expectations in the economy. *Theory and Society* 42 (3): 219–240.
- Beckert, J. 2016. *Imagined futures*. Cambridge, MA: Harvard University Press.
- Benford, R.D., and D.A. Snow. 2000. Framing processes and social movements: An overview and assessment. *Annual Review of Sociology* 26: 611–639.
- Bjerg, O. 2016. How is bitcoin money? *Theory, Culture & Society* 33 (1): 53–72.
- Blanc, J. 2006. *Exclusion et liens financiers: Monnaies sociales, rapport 2005–2006*. Paris: Economica.
- Busco, C., and P. Quattrone. 2018. In search of the “Perfect One”: How accounting as a maieutic machine sustains inventions through generative ‘in-tensions’. *Management Accounting Research* 39: 1–16.
- Cohen, B. 2017. The rise of alternative currencies in post-capitalism. *Journal of Management Studies* 54 (5): 739–746.
- Cohen, B.J. 2004. *The future of money*. Princeton: Princeton University Press.

- Davis, A. 2018. Defining speculative value in the age of financialized capitalism. *The Sociological Review* 66 (1): 3–19.
- De Bakker, F.G., F. Den Hond, B. King, and K. Weber. 2013. Social movements, civil society and corporations: Taking stock and looking ahead. *Organization Studies* 34 (5–6): 573–593.
- Dodd, N. 2016. *The social life of money*. Princeton: Princeton University Press.
- Dodd, N. 2017. Utopian monies: Complementary currencies, Bitcoin, and the social life of money. In *Money talks: Explaining how many really works*, ed. N. Bandelj, F.F. Wherry, and V.A. Zelizer. Princeton: Princeton University Press.
- Dodd, N. 2018. The social life of Bitcoin. *Theory, Culture & Society* 35 (3): 35–56.
- Drucker, P. 1994. *Post-capitalist society*. Harper Business.
- Fabozzi, F.J., S.M. Focardi, and C. Jonas. 2014. *Investment management: A science to teach or an art to learn?* CFA Institute Research Foundation.
- Farias, C. 2017. Money is the root of all evil—Or is it? Recreating culture through everyday neutralizing practices. *Organization Studies* 38 (6): 775–793.
- Feldstein, M. 1998. *The political economy of the European Economic and Monetary Union: Political sources of an economic liability*. National bureau of economic research.
- Fisher, I. 1911. *Purchasing power of money*. New York: Macmillan.
- Fox, J. 2009. *The Myth of the rational market: A history of risk, reward, and delusion on Wall Street*. New York: Harper Business.
- Frenkel, S.H. 1977. *Money, two philosophies: The conflict of trust and authority*. New York: Basil Blackwell.
- Friedland, R., and D.-L. Arjaliès. 2021. Putting things in place: Institutional objects and institutional logics. *Research in the Sociology of Organizations* 71: 45–86.
- Helgesson, C.-F., and F. Muniesa. 2013. For what it's worth: An introduction to valuation studies. *Valuation Studies* 1 (1): 1–10.
- Hsieh, Y.-Y. 2018. *The rise of decentralized autonomous organizations: Coordination and growth within cryptocurrencies*. London, ON: Ivey Business School, Western University.
- Hsieh, Y.-Y., and J.-P. Vergne. 2018. Bitcoin and the rise of decentralized autonomous organizations. *Journal of Organization Design* 7 (1): 14.
- Hutter, M., D. Stark, and A. Berthoin Antal (eds.). 2015. *Moments of valuation: Exploring sites of dissonance*. New York: Oxford University Press.
- Ingham, G. 1996. Money is a social relation. *Review of Social Economy* 54 (4): 507–529.
- Ingham, G. 1998. On the underdevelopment of the “Sociology of Money”. *Acta Sociologica* 41 (1): 3–18.
- Ingham, G. 2013. *The nature of money*. Wiley.
- Jarzabkowski, P., and S. Kaplan. 2015. Strategy tools-in-use: A framework for understanding “technologies of rationality” in practice. *Strategic Management Journal* 36 (4): 537–558.
- Kaelberer, M. 2007. Trust in the Euro: Exploring the governance of a supra-national Currency. *European Societies* 9 (4): 623–642.

- Knapp, G. F. 1924. The state theory of money. *History of Economic Thought Books*.
- Laborde, S. 2019. *Relative theory of money* (Version 2.718). RTM Multi Format (open access) <https://en.trm.creationmonetaire.info/>. Accessed 16 January 2019.
- Leonardi, P.M. 2013. Theoretical foundations for the study of sociomateriality. *Information and Organization* 23 (2): 59–76.
- MacKenzie, D. 2011. The credit crisis as a problem in the sociology of knowledge. *American Journal of Sociology* 116 (6): 1778–1841.
- MacKenzie, D., D. Beunza, Y. Millo, and J.P. Pardo-Guerra. 2012. Drilling through the Allegheny Mountains: Liquidity, materiality and high-frequency trading. *Journal of Cultural Economy* 5 (3): 279–296.
- Maurer, B., T.C. Nelms, and L. Swartz. 2013. “When perhaps the real problem is money itself”: The practical materiality of Bitcoin. *Social Semiotics* 23 (2): 261–277.
- McAdam, D., and R.W. Scott. 2005. Organizations and movements. In *Social movements and organization theory*, ed. G.F. Davis, D. McAdam, R.W. Scott, and M.N. Zald, 4–40. New York: Cambridge University Press.
- McLeay, M., A. Radia, and R. Thomas. 2014. *Money creation in the modern economy*, Q1. Bank’s Monetary Analysis Directorat: Bank of England.
- Meyer, C., and M. Hudon. 2017. Alternative organizations in finance: Commoning in complementary currencies. *Organization* 24 (5): 629–647.
- Miller, P., and M. Power. 2013. Accounting, organizing, and economizing: Connecting accounting research and organization theory. *The Academy of Management Annals* 7 (1): 557–605.
- Muniesa, F. 2007. Market technologies and the pragmatics of prices. *Economy and Society* 36 (3): 377–395.
- Muniesa, F. 2011. A flank movement in the understanding of valuation. *The Sociological Review* 59 (2): 24–38.
- Nakamoto, S. 2008. Bitcoin: A peer-to-peer electronic cash system. <https://Bitcoin.Org/Bitcoin.Pdf>.
- Orlikowski, W.J. 2007. Sociomaterial practices: Exploring technology at work. *Organization Studies* 28 (9): 1435–1448.
- Peacock, M.S. 2014. Complementary currencies: History, theory, prospects. *Local Economy* 29 (6–7): 708–722.
- Piketty, T. 2013. *Le capital au XXIe siècle*. Le Seuil.
- Proudhon, P.-J. 1849. *Proudhon’s solution of the social problem*. New York: Vanguard press.
- Radford, R.A. 1945. The economic organisation of a POW camp. *Economica* 12 (48): 189–201.
- Savoie, D. 2016. *Un projet de monnaie locale à Montréal*. <http://ici.radio-canada.ca/regions/Montreal/2016/07/04/006-monnaie-locale-montreal-economie-sociale.shtml>.
- Simmel, G. 1904. *The philosophy of money*. Routledge (3rd Edition, 2004).
- Swartz, L. 2018. What was Bitcoin, what will it be? The techno-economic imaginaries of a new money technology. *Cultural Studies* 32 (4): 623–650.

- Veneris, A. 2018. *Where do we go from here? Challenges in the era of crypto decentralization*. University of Toronto, Department of Electrical and Computer Engineering and Department of Computer Science. <https://mmpa.utoronto.ca/sites/files/mmpa/public/shared/pdf/%232%20-%20Andreas%20Veneris.pdf>. Accessed 10 January 2019.
- Vergne, J.-P., and G. Swain. 2017. Categorical anarchy in the UK? The British media's classification of bitcoin and the limits of categorization. In *From categories to categorization: Studies in sociology, organizations and strategy at the crossroads (Research in the Sociology of Organizations, Volume 51)*, ed. R. Durand, N. Granqvist, and A. Tyllström, 185–222. Bingley, UK: Emerald Publishing Limited.
- Vidan, G., and V. Lehdonvirta. 2019. Mine the gap: Bitcoin and the maintenance of trustlessness. *New Media & Society* 21 (1): 42–59.
- Wang, S., and J.-P. Vergne. 2017. Buzz factor or innovation potential: What explains cryptocurrencies' returns? *PLoS ONE* 12 (1): e0169556.
- Williamson, O.E. 1991. Comparative economic organization: The analysis of discrete structural alternatives. *Administrative Science Quarterly*, 269–296.
- Yue, L.Q. 2015. Community constraints on the efficacy of elite mobilization: The issuance of currency substitutes during the panic of 1907. *American Journal of Sociology* 120 (6): 1690–1735.
- Yuran, N. 2014. *What money wants: An economy of desire*. Stanford: Stanford University Press.
- Zelizer, V. 2004. Circuits of commerce. In *Self, social structure, and beliefs: Explorations in sociology*, ed. J.C. Alexander, G.T. Marx, and C. Williams, 122–144. Berkeley: University of California Press.
- Zelizer, V.A. 1989. The social meaning of money: “special monies.” *American Journal of Sociology*, 342–377.
- Zelizer, V.A. 2012. How I became a relational economic sociologist and what does that mean? *Politics & Society* 40 (2): 145–174.



Digital Currencies and Central Banks

Gina Pieters

1 Increasing Reliance on Electronic Payment Systems

The global shift to account-based payments began with the widespread introduction of checks and continues today with the introduction of electronic payment networks like Venmo or PayPal (Demirgüç-Kunt et al. 2017). These payments systems allow economic participants to decrease their reliance on physical forms of currency in economic interactions, giving rise to currency digitization. “Currency digitization” does not require the introduction of a digital currency alternative by a Central Bank but does allow us to anticipate some potential impacts of digital currency. Two countries, Sweden and Kenya, have led the way in mass adoption of electronic transfer systems, and the shifts to digitized currency have very different outcomes in the two countries.

G. Pieters (✉)

Department of Economics, University of Chicago, Chicago, IL, USA

e-mail: gcpeters@uchicago.edu

Centre for Alternative Finance, Judge Business School, University of Cambridge, Cambridge, UK

1.1 Case Study: Sweden Trending Cashless

Figure 1 shows the decline in physical currency in Sweden: In 2007 the value of cash in circulation was SEK 109 billion; by 2017 it fell to SEK 57 billion: approximately halving in one decade. The number of banknotes in circulation has similarly undergone a decline. In a survey of the Swedish population conducted by the Swedish Central Bank (the Riksbank) 13% reported using cash for their last purchase, while most individuals (80%) reported they used a debit card.

Some stores now refuse to accept any physical currency. In 2014, 70% of participants reported no problems paying with cash in a shop, a number that fell to 47% in 2018 (The Riksbank 2018). The increase of non-cash payments reduces the use of ATMs or at banks: 20% of the Riksbank survey respondents reported never withdrawing cash from an ATM or a cash desk at a bank. In turn, banks have cut their costs by removing ATMs and shutting down bank branches, which further decreases access to cash and increases the difficulty for companies to deposit any cash earnings (Rogoff 2016). Those not familiar with, interested in, or able to use the technology required for electronic transactions report that it is increasingly difficult to use physical cash to engage with the economy. The Riksbank survey found that while 48% of the aggregate population feel fairly to very positive about the decline in the use of cash and only 27% are fairly to very negative, in rural areas this changes to 35% fairly to very negative with only 32% fairly to very positive.

In addition to individual difficulty accessing an economy's medium of exchange, the Riksbank has become increasingly concerned over the role that payment intermediaries, the entities that provide the electronic payment network, played as all are private companies (Ingves 2018). The Parliamentary Riksbank Committee proposed that all banks should be obliged to handle physical cash: a proposal the Riksbank felt did not go far enough. The

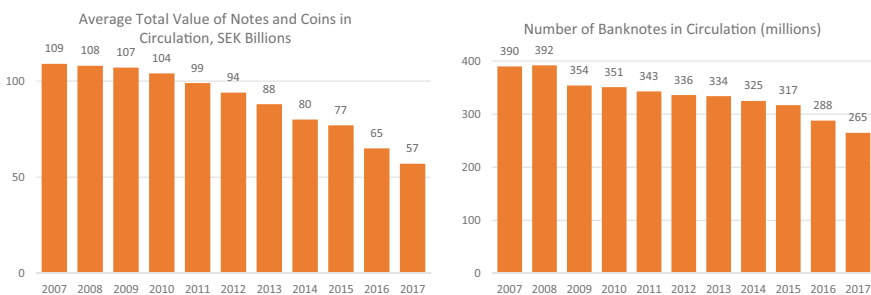


Fig. 1 Sweden has experienced a decline in both the value and number of banknotes and coins in circulation (Source The Riksbank, author calculations)

Riksbank counter proposed that “not only bank, but all other credit institutions that offer payment account should be obligated to handle cash” (Ingves et al. 2018).

1.2 Case Study: Kenya’s m-Pesa

While Sweden has been trying to prevent the collapse of cash to ensure economic access and reduce the role of intermediaries, Kenya’s Central Bank has been embracing currency digitization as an equalizer. In 2007, a major cellphone provider in Kenya created a new program called m-Pesa which allowed consumers to deposit money into their account, either through a link to a bank account or a deposit at a cellphone stall, and transfer it between cellphones. The cellphone company acts like a banking agent, with the m-Pesa system facilitating the electronic transfer of money. Figure 2 shows that the value of transactions using mobile payments has grown sharply since the introduction of m-Pesa: from nearly non-existent in 2007, to over 300 Billion Kenyan shillings today. Unlike Sweden, this increase in currency digitization has not corresponded to a decrease in value of the currency in circulation.

Unlike Sweden’s experience, the increase in non-cash payments has not been accompanied by a decrease of currency in circulation. It has instead increased economic access for the poor, elderly, or those living in rural areas. Prior to the introduction of m-Pesa, many of whom had been previously been unbanked or, because they lived far away from banks or ATMs, had difficulty depositing/withdrawing cash (Jack et al. 2010, 2013). The increase in mobile payments and mobile wallets has allowed these individuals to increase their economic participation.

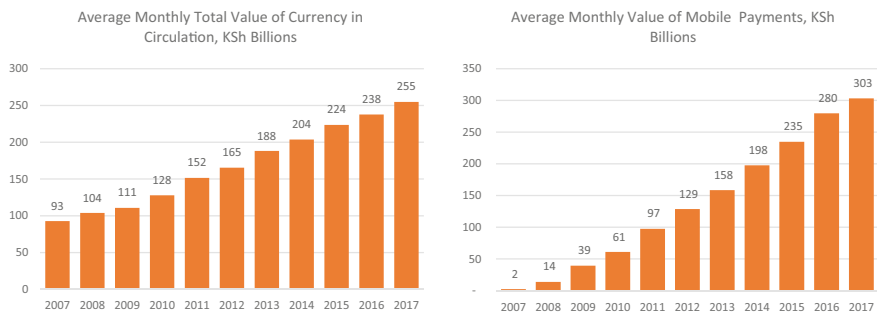


Fig. 2 Mobile payments and currency in circulation for Kenya are both increasing (Source Central Bank of Kenya, authors calculations)

The contrast in these two case studies should serve as a caution to all readers: the impact of a digital currencies on government-issued physical cash may differ drastically from one economy to the next. In the case of central bank issued digital currencies the impact on physical cash may be even harder to disentangle. They simultaneously represent an electronic substitute to physical cash—Sweden’s decrease of cash in the economy—and a means by which to stimulate economic activity and increase transactions—Kenya’s experience.¹

2 Currency Digitization vs. Digital Currency

Sweden and Kenya are instructive regarding the potentially non-uniform impacts of currency digitization, but how does this relate to digital currency? Where do “cryptocurrencies,” “distributed ledger technology,” and “blockchain” fit in? These terms are frequently used interchangeably, but the intersections and distinctions between them are non-trivial:

- *Currency Digitization* refers to the increased use of electronic transfer platforms (which *may* include digital currencies).
- *Digital Currencies* are an electronic representation of money—issued either by a central bank or independent group—that has purchasing power within an economic system. At the most basic technological level, a digital currency requires that users have account(s), and a system that allows users to electronically transfer value between accounts.² The system must also ensure that users cannot double-spend³ funds, or fraudulently transfer funds from another’s account to their own.
- *Distributed Ledger Technology (DLT)* is a class of technology that can be built into a range of electronic systems, including electronic transfer platforms and digital currencies. *Blockchain* is a form of DLT.⁴ Based upon

¹ It should be noted that, in both economies, cash in circulation as a share of GDP has been decreasing. This should not be interpreted as evidence that a digital currency will necessarily encourage economic growth.

² I use “account” to refer to the method by which the storage system identifies a holder to of some amount of digital currency. An account could be as informal as a unique alphanumeric sequence that, when combined with a password, grants the user access to the funds associated with that sequence, or it could be as formal as presenting proof of identity to conform to AML/KYC standards.

³ Double-spending is when the same unit is spent twice. Consider a store issuing coupons: double-spending would occur if a person could create photocopy of a coupon they received and continuously use it.

⁴ Blockchain refers to data that is grouped and sorted, with a consensus mechanism for harmonizing the records across multiple ledgers. This does not inherently require that the consensus mechanism

their protocol design, DLT systems allow for the decentralization of digital currencies. Digital currencies can exist without using a DLT system but are centralized by design if they are.

Both currency digitization and digital currencies require a scalable and robust transfer system for electronic payment systems. A Central Bank could decide to incorporate new technology, such as DLT, to speed up electronic payments (payment system improvements), or it could design a national monetary system in which a physical, state-issued currency is replaced by (or paired with) a digital currency. If a Central Bank does not provide a platform private companies may find it profitable to provide one instead—the role platforms like Venmo or VISA currently occupy.

The reported interest of many Central Banks in Blockchain is not for a digital currency application but rather as a means to improve their payment systems.⁵ Experiments in this vein include Project Jasper from the Bank of Canada (Bank of Canada 2017), Project Ubin by the Monetary Authority of Singapore (Monetary Authority of Singapore 2016), a joint venture between the European Central Bank and the Bank of Japan called Project Stella (European Central Bank 2017), Project Khokha by the South African Reserve Bank (South African Reserve Bank 2018), and Dukaton by the Central Bank of the Netherlands (De Nederlandsche Bank 2018).⁶

Improving efficiency and decreasing technological vulnerability of payments systems has clear economic benefits. The case for Central Bank issued digital currencies is less simple and—as subsequent sections will outline—Central Banks will be impacted by independently issued digital currencies regardless of whether they issue their own. Before advancing to those issues, it will help to develop a shorthand for differentiating key digital currency characteristics.

is decentralized or permissionless. The technology underlying both Blockchain and DLT is discussed in (Rauchs et al. 2018), while an example focused only on Blockchain can be found in (Koch and Pieters 2017).

⁵ VISA has one of the largest centralized electronic payment networks in the world, processing over 24,000 transactions per second. In contrast, Bitcoin processes 7 transactions per second, slower than Ethereum (15 transactions per second) or Ripple (one of the fastest at 1500 transactions per second can process less than 10% of VISA's speed).

⁶ For a full list of announced projects by Central Banks see Appendix A in (Prasad 2018).

3 A Taxonomy of Digital Currencies (DCs)

Two factors are key to establishing a context for the interaction of Central Banks with digital currencies: who has issued the digital currency (independent or government), and whether the digital currency can be thought of as centralized or decentralized (CDC or DDC). In the case of government-issued digital currencies there is also the question of the relationship to government-issued physical currency which will be discussed later.

1. An Independently issued, Centralized Digital Currency (iCDC) (Example: egold, Libra).
2. An Independently issued, Decentralized Digital Currency (iDDC) (Example: Bitcoin).
3. A Government issued, Centralized Digital Currency (gCDC) (Example: China's Digital Yuan project)
 - a. either Equivalent to (equal value, or indistinguishable from) a Central Bank issued physical currency (E-gCDC).
 - b. or in Parallel with a Central Bank issued physical currency, but with some distinguishing market access or prices (P-gCDC).
4. A Government issued, Decentralized Digital Currency (gDDC, with E-gDDC and P-gDDC distinctions).

I use the phrase “Government Issued” instead of “Central Bank issued” to be as broad as possible. I also avoid the term Central Bank Digital Currency (CBDC) as it is unclear with respect to both the de/centralized nature of the currency in question and its relationship to fiat currency. By the taxonomy above, a CBDC may refer to either gCDC or gDDC.

Like other applications of DLT, a digital currency can be a hybrid containing centralization of some components and decentralization of others. The classification of digital currencies into centralized and decentralized can therefore be difficult (Rauchs et al. 2018). For simplicity, this handbook chapter will study only the extremes: digital currencies that are entirely centralized or decentralized, though readers should recognize that any centralized or decentralized features in a hybrid system will inherit the associated benefits or complications highlighted here (Kahn et al. 2018) consider hybrid systems in more depth.

4 Centralized and Decentralized Digital Currencies (CDC vs. DDC)

The distinction between centralized and decentralized digital currencies pivots around who has power to (i) create accounts and (ii) validate value transfers across the electronic payment network. For CDCs a central authority has total control over both aspects, while for DDCs no single entity has an opportunity to deny an account or force a transaction.⁷ A hybrid digital currency grants a degree control over one aspect, but not the other. CDCs—whether issued or controlled by a government (gCDC) or an independent entity (iCDC)—can exist without a blockchain (or more broadly DLT), the merits of which have been debated since account-based digital currency was feasible: (King 1999) and (Woodford 2000), for example, are two decades old.

DLT is a pre-requisite to create DDCs, whether initially issued by a government (gDDC) or issued by an independent issuer unaffiliated with a government (iDDC).⁸ DDCs share many of the same opportunities and obstacles presented by CDCs, but with an additional phenomenon. The decentralization of the currency means that governments, companies, and regulators are unable to shut down a central node to prevent payment transfers or deny access to the DDC.

An example of independent centralized digital currency (iCDC) was eGold. Launched in 1996, issued by a company called Gold & Silver Reserve Inc, accounts on eGold held claims to grams of gold, and could transfer those values between other accounts on eGold. The role of eGold in criminal enterprise, fraud, and hacking attacks led to an investigation of the company, and its subsequent suspension in 2009.⁹ The vulnerability of iCDCs to government investigation and shutdown is why most dismiss the notion of a centralized independent digital currency as a serious contender to a physical state-issued currency.

Bitcoin is an example of an independent decentralized digital currency (iDDC), first described in (Nakamoto 2008). Anyone wanting to use Bitcoin can download the freely distributed program to create their any number

⁷ Centralized power over accounts means that there are scenarios under which an account can be shut down or denied. A decentralized account system exists in a system like Bitcoin, where anyone can download the software and anonymously establish a unique identification number.

⁸ The true identity of the creator of Bitcoin, Satoshi Nakamoto, is famously unknown, and could represent an individual or group of individuals. This is in contrast to companies such as Ripple that have created and maintain their own cryptocurrencies.

⁹ For the story of eGold see (Zetter 2009), while (Foley 2013) discusses the application to Bitcoin exchanges.

of “wallets,” which assigns them an account number. Once miners (who function as auditors or accountants in Bitcoin) “confirm” that a wallet has a sufficient balance for a transaction and incorporate a transfer into the blockchain (the record of all Bitcoin transfers) the transfer is complete. The blockchain (a complete record of all transfers) is copied and distributed to all miners—there is no master copy or master miner that decides whether a transaction should be registered.¹⁰ As long as at least one miner is fully functional somewhere in the world, Bitcoin transactions will be processed, and the record of account balances will continue to exist. The only way for a government actor to undermine the Bitcoin network would be through the acquisition of a computer so powerful that they can “take over” the record writing process, crowding out private actors with superior computing power.¹¹ This robustness to external interference in records is one of the distinguishing features of decentralized digital currencies, and is usually referred to as either tamper-resistance or censorship-resistance.

5 Digital Equivalent or Digital Parallel (E-x or P-x)

A government may choose to issue a digital currency as a one-to-one digital representation of their physical currency with no underlying distinction except for a lack of physical form (E-gDC). An E-gDC has equivalent purchasing power as the equivalent physical issued currency. Alternatively, the government could issue the digital currency as a secondary, parallel currency (P-gDC). A P-gDC functions as a second monetary unit of account and can trade with the physical currency at a non-par exchange rate. For example, if the government issues the P-gDC as a replacement for the physical currency—declaring that the physical currency is no longer acceptable as a medium of exchange—the physical currency would have a government stated value of zero.¹² A P-gDC could also be a currency designed for use exclusively in domestic or international markets, for social programs such

¹⁰ For the purpose of this chapter, unless explicitly noted otherwise, I will treat all blockchains as fully functional and uncompromised.

¹¹ The problem this action faces is that a decentralized DLT-based currency can fork—issuing users an equivalent amount of digital currency on the new network which is also decentralized but not under government control—meaning that any intervention can be bypassed through a new digital currency. This is key to the effective decentralization of a cryptocurrency. A blockchain, or DLT system, is necessary but not sufficient to guarantee decentralization.

¹² The black-market value of the physical currency may differ from the government decreed value.

as food assistance, or in payment for government contracts to government workers.

A P-gDC can be created unintentionally, for example through the design of the associated electronic payment network. Suppose that a consumer wishes to pay a company using a Central Bank-issued digital currency that was intended to act as an E-gDC. If the consumer is required to use the Central Bank's electronic payment system to transfer the gDC, and that system is significantly slower, costlier, or more prone to error than that of private entities, individuals or businesses may choose not to accept the gDC. This would decrease its usefulness as a means of payment. The divisibility of the units of a "E"-gDC to either higher or lower transactional units than the physical currency can also create an unintentional mismatch between a gDC and its physical counterpart.

As previously stated, a necessary requirement for a digital currency is an account to identify the ownership of the balance. One version of a gDC allows consumers the option of creating national savings account directly with the Central Bank. By allowing savings accounts the Central Bank can directly issue currency to consumers in electronic form (for example, through interest rate payments) without printing additional physical currency. This implementation could allow for the co-existence of both physical and electronic currency issued by a Central Bank, as well as an independent electronic payment system by private banks and companies, with payments traveling between various account holders. In this implementation, the E-gDC acts merely as a digital representation of a physical currency, with consumers retaining the freedom to move between the digital and physical representations.

However, suppose that the Central Bank limits the entities allowed to open accounts. For example, the Central Bank could decree that the new gDC accounts can only be issued to "good citizen" individuals with national identification numbers and no criminal records. On one hand, this decree would limit the number of consumers with the gDC, making it less ubiquitous and potentially less desirable than physical currency. On the other hand, using the gDC would indicate that the payee satisfied the "good citizen" requirements imposed by the Central Bank or the government. A "good citizen" may find more options renting an apartment or be granted a discount if they are able to use the gDC to pay their rent, which could make it more desirable than the physical currency.

Account-holder restrictions may introduce other distortions in the payment system. If businesses are not allowed to create accounts, how will

they receive the “good citizen” digital currency? Does the Central Bank automatically convert it into physical currency en-route at no cost? Or would business payment providers be forced to convert at a rate specified by the Central Bank? Alternatively, does the Central Bank impose a second set of standards to identify “good businesses” and allow only those businesses to create an account. A system with all these restrictions is highly unlikely to produce a gDC that is equivalent to the physical currency.

6 Privacy and Anonymity

Physical currency allows for anonymous and private payments in exchange for goods and services; without a receipt there is no record of which parties entered into the transaction, when, or what was exchanged. That is not necessarily true of digital currencies. The nature of electronic recordkeeping requires, at a minimum, two account numbers and acknowledgment of a transfer of currencies between them.

It is a common misconception that cryptocurrencies like Bitcoin allow individuals to digitally transfer money anonymously. The Bitcoin system does not have a “gated on-ramp”: there is no intrinsic identity requirement to create an account within the Bitcoin network, anyone can download the required wallet software and generate an account number (or multiple account numbers). Therefore, while the account numbers engaged in a transaction are known, the true identity of any individual(s) associated with the account are not necessarily recorded.¹³ This represents a coding decision common to decentralized digital currencies: to minimize the amount of control any entity can exert over the system there are no requirements to establish an account.¹⁴

Centralized digital currencies do not need to follow this model, and the designer can impose any requirements they deem necessary. Specifically, a CDC can require that all individuals verify their identity before being granted an account number linked to that identity. If the centralized issuer can also observe transactions records they then have information on all payments between account holders. This does not intrinsically divulge what

¹³ Bitcoin transaction records are contained on the blockchain and are viewable: all records of Bitcoin transactions between accounts are public information. For an example, see blockchain.info.

¹⁴ Transactions that use intermediaries, such as exchanges, may require that users disclose identification information as part of their Anti-Money Laundering (AML)/Know Your Customer (KYC) procedure. Under these circumstances, the activity of the users of the platform can be tracked by the platform. This results in differences in behavior between exchanges that adhere to AML/KYC and those that do not (Pieters and Vivanco 2017).

the currency was used to buy, but the designer could embed a “receipt requirement” requiring that both account holders testify to what was received in exchange for the currency. This testimony could be false, but it would still be recorded and could be used as a form of evidence.

7 Centralized and Decentralized Government-Issued Digital Currencies (gCDC vs. gDDC)

The lack of requirements to open an account and the inability of central authorities to shut down or reverse payments makes a DDC easier to use for criminal enterprise than physical currency, as it allows online, electronic payment for goods in a way that current credit-card-based systems do not allow. A government that issues a fully decentralized gDDC can therefore reasonably expect that it will be used to facilitate crime. For this reason, several have argued that a Central Bank Digital Currency must be a centralized gDC (Bech and Garratt 2017; Berentsen and Schär 2018). However (Kahn et al. 2018), consider a range of alternative implementations that consider both centralized and decentralized schemes.

Centralized issuers can link identities with accounts, and the accompanying transparency of payment networks provides them the ability to monitor all accounts and transactions. This greatly increases the difficulty of payments for criminal activity and tax evasion. Depending on the extent of the network control, a centralized issuer could refuse or reverse any transactions between account holders suspected of illicit activity, or block their ability to access any funds associated with their account.¹⁵ Going farther, it could be incorporated with a Social Credit System—such as the one that China is implementing—to further reward or punish citizens (Liang et al. 2018).

At the same time, this argument can be used to oppose a gCDC to prevent potential abuse of citizens. One can imagine a government that limits the amount of a gCDC that individuals could be spent on “sin” goods (whether soda or alcohol), or what they define as non-patriotic groups (whether terrorist organizations or political opposition groups). The ability to suspend transactions or block access to accounts could also be used to punish political opponents, or to punish individuals who expose corruption or illegal activity.

¹⁵ This argument is examined in (Rogoff 2016).

A possible compromise is to allow the continued existence of physical currency in conjunction with a gCDC, allowing individuals the choice of physical currency or digital currency. This, as discussed above, runs the risk of a parallel economic system: those comfortable revealing their transaction records to the government are rewarded as “good,” and those who are not penalized as “bad.” China’s experiments with social scoring show that this is not merely dystopian fiction; it can result in a bifurcated economy with outsized returns going to government-approved economic or social behavior.

However, if a physical currency continues to exist alongside a gCDC the ability to engage in criminal transactions is unabated which potentially negates a reason to switch to a gCDC in the first place. Suggestions explored in (Rogoff 2016) seek to minimize criminal activity by limiting the maximum denomination of the physical currency, or issue physical currency only in cumbersome physical form—either weight or size—to dissuade its use in large or criminal transactions. Both approaches would negatively impact innocent populations unwilling or incapable of using the technology involved in digital currencies.

Another reason to continue to physical currency issuances is the potential catastrophic failure of a digital currency, whether from hostile action or “acts of god.” A natural disaster or attack that takes out an electric grid, or information transmission systems like the internet, would make it impossible for electronic payments to be processed. In the absence of a physical currency, individuals would be forced to resort to a barter or trust-based system. This would exacerbate economic stress in a time of dire need.

In a payment system disruption, system can function normally once restored; however, there are other, more devastating, cyber-attacks. A cyber-attack could rewrite account balances or cancel and insert transactions, making the underlying transaction or account value data suspect and undermining trust in the system. Such attacks would be the digital equivalent to war-time strategies in which counterfeit currencies are introduced by enemies to generate inflation or sow instability. In a scenario where the government is unable to shift smoothly to a physical currency counterpart, the failure of trust could permanently undermine the entirety of the money supply.

8 Monetary Policy with Centralized Digital Currency (gCDC) Accounts

Most Central Banks limit access to their reserves to select financial institutions that engage in large-value payments system. One way in which a gCDC

could be implemented is to allow individuals to establish their own wallets (the equivalent of their own centralized deposit accounts), with interest rate payments on these accounts equivalent to new issuances of the gCDC. The funds in these accounts could circulate through the economy without additional involvement from the Central Bank beyond the potential provision of the ledger processing, or they could be restricted to only transfers between other Central Bank accounts. This implementation requires centralization of the account system as individuals are opening account with the Central Bank specifically. Processing of payments, on the other hand, *may* be centralized (with funds transferred between accounts and processed by the Central Bank), or *may* be decentralized using DLT, or a hybrid using public-private partnership similar to M-Pesa.¹⁶

A model for Central Bank saving accounts for private individuals is reviewed in (Bordo and Levin 2017). Such a deposit could serve as a haven account during times of financial instability or banking crisis, which might also reduce financial system instability and capital flight, softening the severity of short-run economic fluctuations; it could also drain the financial system of funds for banks to loan to individuals and companies¹⁷ (Barrdear and Kumhof 2016; Keister and Sanches 2018; (Davoodalhosseini 2018). All predict that the gain in payment efficiency, or the gains from increasing bank competition for savers funds, is sufficient to generate positive economic welfare results.

In such a system interest rates on wallets become an additional interest rate policy tool for the Central Bank. The Central Bank could also engage in “helicopter drops” of money: making deposits to all wallet owners as economic stimulus. Of course, the idea of allowing individuals accounts at a Central Bank is not a new notion, nor does it require a digital currency to implement. Implementing a gCDC *does* have the advantage of bypassing the need for physical locations or branches, a substantial barrier to the establishment of physical Central Bank savings accounts (Prasad 2018).

If Central Banks can deny or hinder consumers’ ability to withdraw their funds into alternative savings vehicles (like physical currency), the interest rate on savings accounts at the Central Bank could be negative: a monetary policy tool that is unconstrained by Zero-Lower-Bound concerns.¹⁸ Such a

¹⁶ Variations of this implementation are examined in depth in (Bordo and Levin 2017) and (Davoodalhosseini and Rivadeneyra 2018).

¹⁷ Proposals for Central Bank savings accounts assume that the Central Bank does not extend loans: this is to reduce incentive problems and ensure system stability. This means that the standard financial system retains the responsibility of extending loans to private ventures and individuals.

¹⁸ An example of such a policy could be a fee that is incurred if an account does not engage in a sufficient volume of transactions or purchases each month.

policy also raises the specter of unequal impact across consumers. Consumers with high savings and income may be able to tolerate such negative interest rates and refuse to increase their spending, while consumers with low income or wealth are forced to spend. Alternatively, consumers with high balances may be able to easily spend the required amount of money to avoid a fee, while consumers with low balances may be unable to do so. Both forces would increase economic inequality and would presumably do so at the worst time for marginal populations, when the Central Bank is attempting to use negative interest rates—a recession. This effect further compounds with the heterogenous impact that digital currencies may have on different segments of the population, as noted above in the currency digitization case studies of Sweden and Kenya.

9 Programmable Money and Stablecoins

Stablecoins are a category of digital currency that maintain a stable exchange rate to a chosen target.¹⁹ Traditionally there are two ways that such an exchange rate pegs are maintained. One is through the accumulation of the target good, which can be redeemed upon receipt of currency. The cryptocurrency Tether claims to use this reserve-based method to maintain a peg of 1 Tether to US\$1: all tethers are redeemable for US\$1. Many financial regulations on markets using the US dollar do not apply if only cryptocurrencies are used—therefore the use of Tether as a non-US dollar dollar-equivalent has made it a very popular cryptocurrency despite concerns about the veracity of their claim (Griffin and Shams 2018).

An alternative method to maintain a peg is through monetary policy: issuing, destroying, or exchanging money to change its exchange rate relative to the target; in Central Banks this is managed through the monetary policy committee. A CDC could use the same system. However, a DDC—with no centralized guiding committee—could only use a monetary policy which has been hard coded into it—programmable money. Bitcoin issuances of new bitcoin at a fixed rate to miners is an example of programmed monetary policy, albeit one that does not respond to Bitcoin prices or the Bitcoin economy.

There are several iDC attempts to issue stable coins with encoded monetary policy to maintain exchange rate stability. At this time, efforts to achieve such encoding in a decentralized system are still stymied by technological

¹⁹ Moin et al. (2019) introduce a classification framework for stablecoins based on their design elements.

realities: using data such as prices or exchange rates relies on an outside data source. Most blockchains track only the quantity of coins moved between the account, not what has been received in return, and therefore price data is not native.²⁰ The quantity of coins traded cannot be used as a policy variable for monetary policy without gate-keeping accounts as there is no guarantee that they belong to unique individuals. This allows individuals to engage in self-trading, artificially increasing the trade volume of a cryptocurrency. These two factors make it difficult to efficiently use a purely internal, contained data source to enact monetary policy.

A reliance upon external data exposes the cryptocurrencies blockchain to a weakness: With exogenous references a system is no longer self-contained and self-directing. External data can be falsely reported, manipulated, or unavailable, any of which could cause the system to enact incorrect policies (Rauchs et al. 2018). Consider, for example, the range of data monetary policymakers consider in making decisions: interest rates, unemployment rates, job vacancies, turnovers, quits, inventory data, shipping data, even recent weather events—domestically and with major trade partners. A monetary policy-based stablecoin must therefore either rely on a vast network of exogenous references—each introducing its own potential source of manipulation—or be simplistic and rely on purely internal statistics which may themselves be manipulatable.

10 Money Competition

Bitcoin was initially intended to be used as an alternative to Central Bank issued currency, though at this time there is no evidence that any iDC is used to purchase goods on a large scale in any economy (Jonker 2018). Its primary use seems to be as a speculative asset (Baek and Elbeck 2015). While this may seem natural, it is worthwhile noting that having an independently issued currency is not unheard of historically.

Should the use of iDCs to purchase goods increase, cryptocurrencies represent a return to a time in which banks could issue their own bank notes, with a network of acceptances. This era came to an end in most countries with the establishment of Central Banks as the sole issuing authority for currency.²¹

²⁰ The Bitcoin blockchain, for example, can record that you sent 1 bitcoin to someone, it cannot record that you received US\$10,000 in return. Both sides of the transactions are needed to deduce that the price of 1 bitcoin is \$10,000.

²¹ A brief summary of the history of this period in the United States can be found in Federal Reserve Bank of Philadelphia (2016) and White et al. (2015).

During the time of bank issued currencies there was concern about financial system fragility: the collapse of a bank could lead to a collapse in the bank notes associated with that bank, in turn wiping out the savings and purchasing power of any associated bank note holders. While such shocks were rare, it is not obvious that there are analogous concerns with a fully decentralized cryptocurrency which is not associated with a central issuing company.

The launch of novelty customer loyalty “WhopperCoin” by a Burger King in Russia raises a different application (BBC News 2017). WhopperCoin is a marketing variation of a standard consumer loyalty program. However, companies could try to issue iCDC as a variation of a company scrip. In the nineteenth century, this practice was common in remote regions where companies found cash difficult to obtain. Workers would collect their wages not in currency, but in scrips, but could then redeem the scrips only in company stores where goods were sold at large mark-ups due to the company’s monopoly power. If instead company scrips were issued as iCDCs that could be spent in any store that accepts them then this monopoly exploitation may no longer be of concern. The emergence of online companies that can easily convert prices denominated in various currencies would facilitate this.

In addition to concerns about privacy already outlined primarily in the context of a government issuer, this also raises the question of whether the company-issuer has the right to control on-ramping or off-ramping. For example, if McDonalds issues and pays workers in a cryptocurrency (McCoin), would it allow its competitor BurgerKing to accept and process transactions paid for in McCoin’s? Would Burger King or McDonald be allowed to set a higher or lower price for individuals using McCoin’s? Such questions are currently in the realm of speculation.

11 Decentralized Digital Currencies (DDCs) in the Global Economy

Decentralized digital currencies have no single entity that approves or processes transactions or users, and therefore no single entity that regulators can target to shut down the entire ecosystem. DDCs are already being used to avoid capital controls and exchange rate manipulation in countries that employ such policies. This has two implications for Central Banks: firstly, DDCs undermine any attempt to use capital controls or exchange rates in

monetary policy; subsequently, DDCs have the potential to create a new era of harmonized interest rates and prices.

A country that wishes to control their exchange rate must either give up monetary policy or use capital controls (Pieters 2017). Most countries that have a managed exchange rate use capital controls, allowing use of monetary policy to reduce economic fluctuations. DDCs allow the evasion of capital controls, removing this degree of policy freedom and forcing Central Banks to choose between managed exchange rate and monetary policy freedom (Pieters 2016) (Fig. 3).

While a government may attempt to force a certain rate of conversion (or a fixed exchange rate) to a specific digital currency, there is no evidence that they will be able to do so given the global trading of DDCs. In this way, DDCs represent an analogy to a return to the era of the gold standard: if all state-issued currencies can be freely converted into Bitcoin (or any other DDC) any differences between monetary systems must collapse through arbitrage seeking in both the financial and traded goods markets. Unlike gold, however, a DDC cannot be blockaded or confiscated.

This poses an additional constraint upon Central Bank policies as the ability to independently set interest rates will be undermined by a globally available alternative. The global linkage of assets, interest rates and exchange

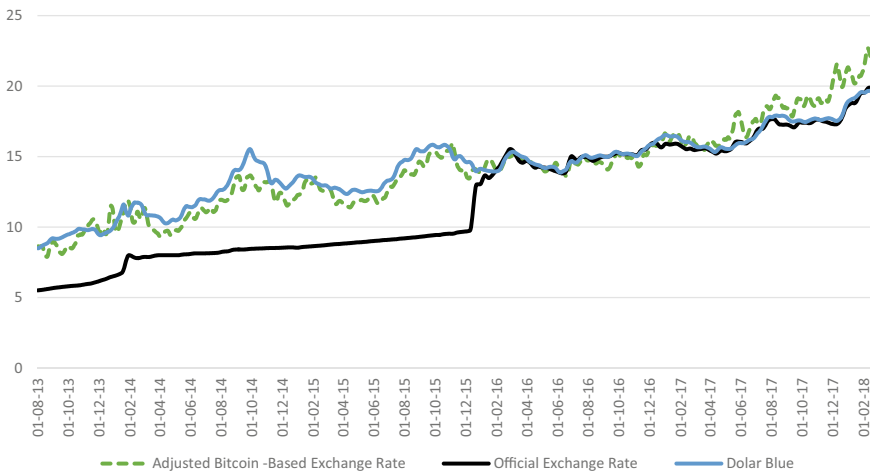


Fig. 3 Relative Bitcoin prices move in a way that reflect the unofficial exchange rate, regardless of official exchange rate policy. This implies that individuals are using Bitcoin to circumvent capital controls and exchange rate pegs. Data Argentinian Peso to US dollar Exchange Rate: Official, Street, and Bitcoin-based. Peg in the 2013–2015 I period, moving to floating exchange rate in 2016 (Source LocalBitcoins pricing data, Updated authors calculations based on Pieters [2016])

rates would in turn link inflation rates around the world. In the same way that adopting the Euro has created convergence of prices throughout the European Monetary Union, a “Bitcoin Standard” would create a convergence of inflation rates throughout the world, *if* consumers use the digital currency for purchasing goods (Weber 2016; Benigno et al. 2019).

Astute readers will notice that I have confined the discussion here to DDCs. CDCs must have a nexus: for example, the Libra will be controlled by a Foundation. A government could force Facebook, or any member of the foundation, to deny services within the borders of a country, or prevent acceptance of a certain currency, which would weaken the ability of that CDC to disrupt pre-existing monetary policy. It is more likely to improve or increase international payment transactions or remittance flows along pre-existing networks.

12 Case Study: The Venezuelan Petro

Venezuela has been a first mover in digital currencies. During a period of hyperinflation of the Venezuelan Bolivar, president Nicolas Maduro announced that the country would adopt a cryptocurrency (DDC)—the Petro—which would be pegged to barrels of gasoline: a commodity cryptocurrency. There are no salient details on how the peg is maintained, and the Venezuelan congress has called the Petro illegal.²² Despite this, Venezuela plans to present the Petro to OPEC in a bid to become to digital currency of oil transactions.²³

Venezuela was able to move quickly in launching this initiative because they did not design their own blockchain. Instead Venezuela used a pre-existing blockchain, NEM, issuing what should strictly be called a cryptotoken. Because of the Petro’s architecture neither the Central Bank nor the government has control over the Petro’s blockchain, therefore it is not strictly obvious how the Petro’s status should be considered vis-à-vis the Venezuelan Central Bank.

Subsequent announcements have continued to advance the Petro as the preferred unit of account for Venezuela, with government pension payments issued in Petro, replacing issuances of the Bolivar. Individuals wishing to buy passports from the Venezuelan government must do so using Petros. As in

²² The whitepaper for the Petro can be found here: <https://whitepaperdatabase.com/venezuela-petro-cryptocurrency-ptr-english-whitepaper/>.

²³ Announced in a tweet by the minister of petroleum and the president of the state-owned oil company <https://www.coindesk.com/venezuela-to-present-petro-at-opec-as-the-digital-currency-for-oil/>.

other countries, however, the use of DDCs as a purchasing currency has been minimal, with little evidence the Petro is in wide circulation, and no evidence that the oil which supposedly backs the Petro even exists (Ellsworth 2018).

While the anonymity and decentralization of blockchain-based cryptocurrency may be a detriment in the eyes of many Central Banks, the Venezuelan experience shows how it can be considered to benefit a government. Venezuela has been placed under numerous international sanctions limiting its ability as a major oil exporter to engage in international markets, hastening its economic collapse. As detailed in the Petro whitepaper, one reason Maduro opted to issue a DDC, instead of redenominating the Bolivar, is to circumvent these sanctions in hopes that this would allow Venezuela access to the funds that otherwise have trouble entering the economy.

DDCs may be used in other ways in Venezuela. As the Bolivar undergoes hyperinflation, there have been reports that citizens have turned to mining as an alternate source of income, and (non-Petro) DDCs as an alternate source of savings. Mining a DDC yields that DDC as payment: the DDC can then be converted into a different currency of the miner's choice. While DDC prices are wildly volatile, in an economy experiencing hyperinflation and economic collapse, such as Venezuela's, they may represent a more stable alternative than the state-issued fiat currency earned working normal jobs.²⁴

13 Conclusion on Central Banks and Digital Currency

The numerous tensions surrounding digital currencies are explored in this chapter. If the Central Bank issues digital currency simultaneously with a physical currency it then runs the risk of created a parallel economic system, with no benefit from a decrease in criminal activity or an expansion of monetary policy options to include negative interest rates. However, ceasing physical currency makes the economy reliant on electronic systems that may be vulnerable to cyber-attack or infrastructural failure, and excluding sections of the population that are unable or unwilling to engage with digital currency. If Central Banks opt to issue gDCs, it is unlikely that they will be completely decentralized, however, this raises substantial privacy concerns associated with access to, control over, and responsibility for immense amounts of personal data and freedom of the digital currency users.

²⁴ Many stories circulate on online forums, such as Reddit or Twitter. A more formal examination can be found in (A.F. 2018).

The experience of Kenya and Sweden in relation to currency digitization provide evidence that different economies will have different experiences with digital currencies, whether government issued or independent. The example of the Petro proves this. Even though most Central Banks are not interested in gDDCs, Venezuela issued the Petro specifically to have a means to allow circumvention of international barriers—to encourage what other countries would define as criminal activity—for the benefit of its own economy.

Whether or not a Central Bank issues a digital currency, independent entities are already issuing their own. In the case of iCDCs regulations or restrictions can be levied to reduce monetary competition but DDCs have no central nexus on which to apply leverage. Ultimately, the money competition from DDCs, the decreasing the use of physical currency, and the potential for new global economy linkages mean that Central Banks cannot explore issues related to digital currency entirely on their own schedule.

References

- A.F. 2018. Why are Venezuelans mining so much bitcoin? *The Economist*.
- Baek, C., and M. Elbeck. 2015. Bitcoins as an investment or speculative vehicle? A first look. *Applied Economics Letters*, 30–34.
- Bank of Canada. 2017. *Project Jasper*. Payments Canada. February 9. <https://www.payments.ca/industry-info/our-research/project-jasper>.
- Barrdear, John, and Michael Kumhof. 2016. The macroeconomics of central bank issued digital currencies. *Bank of England Staff Working Papers* 605.
- BBC News. 2017. Burger King launches WhopperCoin crypto-cash in Russia. *BBC News—Technology*, August 29. <https://www.bbc.com/news/technology-41082388>.
- Bech, Morton, and Rodney Garratt. 2017. Central bank cryptocurrencies. *BIS Quarterly Review*, 55–70.
- Benigno, Pierpaolo, Linda M. Schilling, and Harald Uhlig. 2019. Cryptocurrencies, currency competition, and the impossible trinity. *Benigno, pierpaolo and schilling, linda and ublig, harald, cryptocurrencies, currency competition, and the impossible trinity (August 19, 2019)*. University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2019-108.
- Berentsen, Aleksander, and Fabian Schär. 2018. The case for central bank electronic money and the non-case for central bank cryptocurrencies. *Federal Reserve Bank of St. Louis Review*, Second Quarter.
- Bordo, Michael D., and Andrew T. Levin. 2017. Central bank digital currency and the future of monetary policy. *NBER Working Paper* 23711.
- Davoodahosseini, Mohammad, and Francisco Rivadeneyra. 2018. A policy framework for e-money: A report on bank of canada research. *Staff Discussion Paper*, April.

- Davoodalhosseini, M. 2018. Central Bank digital currency and monetary policy. *Bank of Canada Working Papers (2018-36)*.
- De Nederlandsche Bank. 2018. *DNBulletin: DNB is experimenting with blockchain*. June 7. <https://www.dnb.nl/nieuws/nieuwsoverzicht-en-archieef/DNBulletin2018/dnb376502.jsp#>.
- Demirgüç-Kunt, Asli, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess. 2017. *The global finex database: Measuring financial inclusion and the fintech revolution*. World Bank Group.
- Ellsworth, Brian. 2018. Special report: In Venezuela, new cryptocurrency is nowhere to be found. *Reuters*, August 30. <https://www.reuters.com/article/us-cryptocurrency-venezuela-specialreport/special-report-in-venezuela-new-cryptocurrency-is-nowhere-to-be-found-idUSKCN1LF15U>.
- European Central Bank. 2017. *Exploring fintech*, September. https://www.ecb.europa.eu/paym/intro/mip-online/2017/html/201709_article_stella.en.html.
- Federal Reserve Bank of Philadelphia. 2016. The state and national banking eras: A chapter in the history of central banking. *Education Series*, December. <https://www.philadelphiafed.org/-/media/publications/economic-education/state-and-national-banking-eras.pdf>.
- Foley, Stephen. 2013. Bitcoin needs to learn from past e-currency failures. *Financial Times*, November 28.
- Griffin, John M., and Amin Shams. 2018. Is bitcoin really un-tethered? *SSRN Working Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3195066.
- Ingves, Stefan. 2018. Payments in the future and legal protection for the Swedish Krona (DN Debatt). *Sveriges Riksbank*, February 27. <https://www.riksbank.se/en-gb/press-and-published/notices-and-press-releases/debate-articles/payments-in-the-future-and-legal-protection-for-the-swedish-krona/>.
- Ingves, Stefan, Kerstin af Jochnick, Per Jansson, Martin Floden, Henry Ohlsson, and Cecilia Skingsley. 2018. Consultation response on secure access to cash, October 17. <https://www.riksbank.se/globalassets/media/remisser/riksbankens-remissvar/engelska/2018/consultation-response-on-secure-access-to-cash.pdf>.
- Jack, William, Adam Ray, and Tavneet Suri. 2013. Transaction networks: Evidence from mobile money in Kenya. *American Economic Review*, 356–361.
- Jack, William, Tavneet Suri, and Robert Townsend. 2010. Monetary theory and electronic money: Reflections on the Kenyan experience. *Economic Quarterly (Econ)*, 83–122.
- Jonker, Nicole. 2018. What drives bitcoin adoption by retailers. *DNB Working Paper 585*, February.
- Kahn, Charles M., Francisco Rivadeneyra, and Tsz-Nga Wong. 2018. Should the central bank issue e-money? Available at SSRN: <https://ssrn.com/abstract=3271654>.
- Keister, Todd, and Daniel Sanches. 2018. Should central banks issue digital currency? *Working Paper*.

- King, Mervyn. 1999. Challenges for monetary policy: New and old. *Speech delivered at a symposium by the Federal Reserve Bank of Kansas City, Jackson Hole, WY*, August 27.
- Koch, Christoffer, and Gina C. Pieters. 2017. Blockchain technology disrupting traditional record systems. *Financial Insights* 6 (2): 1–4.
- Liang, Fan, Vishnupriya Das, Nadiya Kostyuk, and Muzammil M. Hussain. 2018. Constructing a data-driven society: China's social credit system as a state surveillance infrastructure. *Policy & Internet*.
- Moin, Amani, Emin Gün Sirer, and Kevin Sekniqi. 2019. A classification framework for stablecoin designs. [arXiv:1910.10098](https://arxiv.org/abs/1910.10098).
- Monetary Authority of Singapore. 2016. Singapore Financial Centre, November 16. <http://www.mas.gov.sg/Singapore-Financial-Centre/Smart-Financial-Centre/Project-Ubin.aspx>.
- Nakamoto, Satoshi. 2008. Bitcoin: A peer-to-peer electronic cash system, October. www.bitcoin.org.
- Pieters, Gina. 2016. Does Bitcoin reveal new information about exchange rates and capital controls? *Globalization and Monetary Policy Institute Working Paper 292*, Federal Reserve Bank of Dallas.
- Pieters, Gina. 2017. The potential impact of decentralized virtual currencies on monetary policy. *Federal Reserve Bank of Dallas Globalization and Monetary Policy Institute 2016 Annual Research Report*. <https://www.dallasfed.org/-/media/documents/institute/annual/2016/annual16d.pdf>.
- Pieters, Gina, and Sofia Vivanco. 2017. Financial Regulations and price inconsistencies across Bitcoin Markets. *Information Economics and Policy* 39: 1–14.
- Prasad, Eswar. 2018. Central banking in a digital age: Stock-taking and preliminary thoughts. *Brookings Institution Working Paper*.
- Rauchs, Michel, Andrew Glidden, Brian Gordon, Gina Pieters, Martin Recanatini, François Rostand, Kathryn Vagneur, and Bryan Zheng Zhang. 2018. *Distributed ledger technology systems: A conceptual framework*. Cambridge: University of Cambridge, Cambridge Centre for Alternative Finance.
- Rogoff, Ken. 2016. *The curse of cash*. Princeton: University Press.
- South African Reserve Bank. 2018. Project Khokha: Exploring the use of distributed ledger technology for interbank payments settlement in South Africa.
- The Riksbank. 2018. *Payment statistics*. <https://www.riksbank.se/en-gb/statistics/payments-notes-and-coins/payment-statistics/>.
- Weber, Warren. 2016. A bitcoin standard: Lessons from the gold standard. *Bank of Canada Staff Working Paper 2016-14*.
- White, Lawrence H., Viktor Vanberg, and Ekkehard Köhler. 2015. Free banking in history and theory. In *Renewing the search for a monetary constitution*. Washington, DC: Cato Institute.
- Woodford, Michael. 2000. Monetary policy in a world without money. *International Finance* 3 (2): 229–260.
- Zetter, Kim. 2009. Bullion and bandits: The improbable rise and fall of e-gold. *WIRED*, June 9.



Cryptocurrencies: Regulatory Perspectives and Implications for Investors

Julapa Jagtiani, Michael Papaioannou, George Tsetsekos,
Erik Dolson, and Dan Milo

1 Introduction

Cryptocurrencies are part of the latest innovations in the financial system. They may be defined broadly as digital representations of value issued by private developers and denominated in their own unit of account. Cryptocurrency exchanges operate as “gatekeepers” that serve as a bridge between cryptocurrency schemes and the wider economy. On a blockchain platform,

Disclaimer: The opinions expressed in this article are the authors’ own views and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System, or the International Monetary Fund, its Executive Board, or IMF Management. Any errors or omissions are the responsibility of the authors.

J. Jagtiani (✉) · E. Dolson
Federal Reserve Bank of Philadelphia, Philadelphia, PA, USA
e-mail: julapa.jagtiani@phil.frb.org

M. Papaioannou
International Monetary Fund, Washington, DC, USA

G. Tsetsekos
Drexel University, Philadelphia, PA, USA
e-mail: tsetsekos@drexel.edu

D. Milo
Rutgers University, New Brunswick, NJ, USA
e-mail: dan.milo@rutgers.edu

cryptocurrencies could potentially allow for instantaneous processing and validation of payments and other transactions between parties, without the need for a trusted third party (such as a bank or central bank). In addition, the process would be almost completely anonymous. The potential for a blockchain-based platform and cryptocurrencies to drastically increase efficiency in payment services, securities settlements, trade finance, and other areas of finance, has been a main reason for many analysts to propose that blockchain technology could become a mainstream financial technology for the future.

There are 4136 cryptocurrencies in existence as of May 2020 (CoinLore 2020). This number does not include many that have failed. The most prominent examples of cryptocurrencies are Bitcoin and Ethereum, which have been traded on the Bitcoin and Ethereum blockchain platforms, respectively. Not only have cryptocurrencies become a popular topic of conversation around the globe, but several research studies have recently explored the challenges and opportunities of potential disruption to central bank operations, the payment clearing and settlement systems, as well as the associated risk to financial stability, the financial market, and overall system.

In this context, several interesting questions have been raised—whether cryptocurrencies could replace official central bank currencies, what a cashless financial system would look like, how monetary policy would need to adjust to remain effective, what type of regulatory changes would be needed to close the loopholes around capital controls and effective prevention of anti-money laundering, if central banks should create their own digital currencies, and if central banks should regulate them. There has also been discussion related to the usage of distributed ledger technology (DLT), such as blockchain, in recording and transacting cryptocurrencies and the impact of initial coin offerings (ICOs), which have become a way for start-ups to raise funds.

This chapter is an introduction to cryptocurrencies, their impact on consumers and the financial system, and their regulations.¹ We first briefly discuss the mechanics of how cryptocurrencies work, highlighting the major innovations underpinning their success (i.e., blockchain and other DLT). We also discuss the scope of cryptocurrencies in the global financial system and discuss the emergence of DLT and other fintech innovations and their impact on monetary policy. We present a broad overview of recent trends in the use of cryptocurrencies and will evaluate the potential impact of cryptocurrencies on the wider financial system, and we evaluate the potential of creating a central bank digital currency (CBDC). We also explore regulatory options

¹ For a review of the roles of cryptocurrencies in the global economy, see Jagtiani et al. (2019).

and the important considerations for regulators and policymakers as they attempt to develop an innovation-friendly regulatory system with the goals of finding a balance between promoting innovations and protecting consumers and the financial systems overall.

The rest of the chapter is organized as follows: Section 2 examines the benefits, risks, and regulations of private digital currencies, especially cryptocurrencies. Section 3 discusses central bank digital currencies. Sections 4 and 5 outline the opportunities and challenges of ICOs and the implications for investors, respectively. Section 6 offers concluding remarks.

2 Private Digital Currencies—Cryptocurrencies and Stablecoins

At the 2018 Blockchain Africa Conference in Johannesburg, South Africa, financial experts concurred that cryptocurrencies will boost financial inclusion throughout Africa since trust is a big issue holding banking inclusion back (Mbogo 2018). Private digital currencies also present unique opportunities for less financially developed nations to increase financial inclusion. In many countries, the barriers to credit access are high because of the high cost of finance, limited bank presence, and the lack of trust in the banking sector. If aptly regulated, digital currencies could widen access to banking and financial services by providing a secure, trusted, and transparent way of saving and paying. Beyond this, in unstable countries, cryptocurrencies can serve as a relatively more stable asset than government-issued currencies. One notable example is Venezuela where high inflation (in the hundreds of percent) has led to a significant increase in the use of Bitcoin (Forbes Finance Council 2017).

Despite the potential benefits of cryptocurrencies, critics have raised multiple concerns: their high-price volatility and possible proliferation of fraud, terrorist financing risks, financing of other illicit activities, promoting tax evasion, and financial stability considerations, ultimately arguing for the need to regulate their use. We present some benefits, risks, and regulations relating to private digital currencies, specifically cryptocurrencies.

High Volatility: Some have treated Bitcoin and other digital currencies as an investment vehicle rather than a means of exchange or unit of account such as traditional currency. However, cryptocurrencies have proven to be very risky investments and their prices have been extremely volatile. Some view the Bitcoin price volatility as the biggest bubble in financial history,

while others may consider it more broadly as a short stop on the march to an anonymous and free financial system. The benefits and risks to the users of Bitcoin and other cryptocurrencies and the financial system as a whole are widely debated among policymakers, international financial institutions, and other market participants, while various regulatory proposals are currently deliberated.

Bitcoin saw a dramatic increase in its value from less than \$1 in 2010 to a peak of more than \$19,500 in December 2017, before falling sharply to less than \$4000 in December 2018, and continuing its wide fluctuations thereafter—recently rising to over \$63,000 in April 2021 after the previous decline to around \$5200 in March 2020 (when 50 COVID-19-related deaths were confirmed in the United States). Bitcoin prices were often driven up partly by the expectations of the introduction of Bitcoin future contracts by major exchanges (CBOE Global Markets and CME) in December 2017, coupled with the specified upper limit of 21 million coins in its supply (Hale et al. 2018). The price was then driven downward dramatically by several events—China and Korea moved forward with tighter cryptocurrency regulations in early 2018, the U.S. Securities and Exchange Commission (SEC) imposed new regulations on cryptocurrency trading, and the Bank of England also called for tighter cryptocurrency regulations (Kharpal 2018). These actions, along with the move of major technology firms, such as Google, Facebook, and Twitter, to ban cryptocurrency advertising in the first quarter of 2018, raised investors' concerns that led to a sharp contraction in Bitcoin's demand and a consequent drastic decline in price. Similar price swings also occurred previously—in 2011 and twice in 2013—all attributed to hacking incidences (CoinDesk 2014). Bitcoin has recently started trading and being accepted for online purchases, including on Square and PayPal, driving price up in early 2021.

Illegal Activities: De Filippi (2014) has documented concerns around cryptocurrencies that are being used to fund illegal activities because of their high level of anonymity. Bitcoin has been used in many cases to finance online gambling, black market operations, and other illegal activities. One of the high-profile cases was about a “hidden” website in which users were able to buy and sell drugs, with the majority of transactions occurring through Bitcoin because of the overall untraceable nature of transactions on its network. The Federal Bureau of Investigation (FBI) eventually shut down the site and filed criminal charges against its founder (Ball et al. 2013). This example simply illustrates the potential dangers of highly anonymous digital currencies to the integrity of the financial system.

Impact on Financial Stability: The global hype about cryptocurrencies, especially Bitcoin, and blockchain, has also raised monetary and financial market stability concerns to central banks and financial markets regulators. Quarles (2017) argues that the existence of an alternative asset such as Bitcoin during times of economic adversity may frustrate efforts to stem price and credit risk because the exchange rate for the U.S. dollar may become unstable. Also, a rise in cryptocurrencies could make it difficult for central banks to implement monetary policy since they operate largely through the banking system. Whatever the impact of cryptocurrencies, it is widely believed that the underlying technology that makes cryptocurrencies possible (such as blockchain) will have widespread effects on the financial system.

There have been many different viewpoints regarding the extent to which cryptocurrencies will impact the financial system, ranging from far-reaching to limited. The depth and scope of the impact will largely depend on the extent to which cryptocurrencies are widely adopted, as either money or investment vehicles. Viewpoints about this differ significantly; see Quarles (2017), Bordo and Levin (2017), Raskin and Yermack (2016), and Ali et al. (2014).

To date, cryptocurrency transactions account for only a small fraction of the entire payment system. Cryptocurrencies are accepted as a mean of payment in very few jurisdictions, such as in the Swiss Cantons of Zug (a municipality in Ticino), the Isle of Man, and Mexico (Irrera and Neghaiwi 2018). The number of daily transactions for Bitcoins peaked at about 500,000 per day on average compared with over 300 million average daily transactions processed by the U.S. payments system on a day in 2016 (Blockchain.com 2018; Federal Reserve System 2017). Similarly, a report by the Bank of England argues that digital currencies are not likely to experience wide adoption and that the impact on the financial system will likely be minimal (Ali et al. 2014).²

Similarly, S&P Global Ratings (2018) finds that a collapse in the value of cryptocurrencies would just cause a ripple across the financial services

² They note that the key attraction of digital currencies for users and merchants is the low transaction fees, which are often significantly lower than those for traditional payments systems (credit and debit cards). These low transaction fees are mainly a result of the implicit subsidy given to miners when they receive new currency for verifying transactions. Ali et al. (2014) argue that this subsidy cannot continue indefinitely because of the fixed supply of most digital currencies. As such, digital currencies will eventually have to compete with traditional payment methods on the basis of true cost. Digital currencies tend to have a higher marginal cost than traditional payment methods because of their decentralized nature and increasing returns to scale in the operation of computer servers; thus, digital currencies will eventually become more expensive to produce than traditional currencies, the low transaction fees will rise, and users will no longer use them extensively. Their analysis then implies that widespread adoption of cryptocurrencies is not feasible so their overall effect on the financial system will be quite small even in the long run.

industry and would not disturb financial stability or affect banks' safety and soundness. It also finds that retail investors would be the first to bear the brunt of a collapse in cryptocurrency values. Central bankers have also been concerned about how the transfer mechanism of cryptocurrencies would impact payment system risks. According to Randal K. Quarles (2017), Federal Reserve vice chairman for Supervision:

While these (private) digital currencies may not pose major concerns at their current levels of use, more serious financial stability issues may result if they achieve wide-scale usage. Risk management can act as a mitigant, but if the central asset in a payment system cannot be predictably redeemed for the U.S. dollar at a stable exchange rate in times of adversity, the resulting price risk and potential liquidity and credit risk pose a large challenge for the system.

Regulations: Because the financial ecosystem is changing quickly, the responsibilities of regulation become more acute and complex. The regulatory system needs to remain strict enough to discourage potentially destabilizing market behavior/activities but also flexible and open minded enough to encourage financial evolution and fintech innovations. Cryptocurrencies could potentially broaden access to financial services through mobile banking and provide new channels for financing through crowdfunding. Further, technological advancement such as blockchain, which many financial institutions are currently considering adopting, are undoubtedly destined to drastically alter the existing financial infrastructure and the entire financial landscapes. These financial technology advances pose significant challenges for regulators.

How should central banks and regulators respond to potential opportunities and challenges arising from cryptocurrencies? What measures can be used to curb speculation and what are other main regulatory and legal issues? Market participants are watching carefully where regulations are headed. The ultimate goals would be for cryptocurrency regulations to ensure that the potential benefits of the technology in the financial sector can be realized while the risks are mitigated.

Cross-Border Regulations: Bitcoin and other cryptocurrency transactions are currently subject to different restrictions across countries.³ Regulatory arbitrage may be a concern as regulations tend to vary across jurisdictions.

³ The Law Library of Congress research has explored cryptocurrency regulations around the world, dividing nations into one of these categories: (1) countries that ban all cryptocurrency activity, (2) countries that allow their citizens to use cryptocurrencies only outside of their borders, (3) countries that ban financial institutions from using cryptocurrencies, and (4) countries that believe that the size of the cryptocurrency market is too small to warrant regulation (Global Legal Research Directorate 2018).

Research from the Law Library of Congress has documented international discrepancy on tax policies for cryptocurrencies, e.g., Israel taxing cryptocurrency as an asset, while Switzerland taxing cryptocurrency as a foreign currency (Global Legal Research Directorate 2018). Regulators around the globe have yet to develop a coherent regulatory framework for cryptocurrencies. It appears that regulators still need additional information to feel comfortable to finalize the rules.⁴ The G20 agreed during its meeting in December 2018 to start regulating cryptocurrencies (and crypto assets) for anti-money laundering and countering the financing of terrorism (Suberg 2018).

Product Classification: One of the key issues in the regulation of cryptocurrencies is whether to treat them as currencies, commodities, or securities, each with different, important tax implications. Few countries have considered digital currencies to be actual currencies because research has shown that they do not fulfill the roles of money (Ali et al. 2014). One notable exception is Switzerland, which taxes digital currencies as foreign currencies. Some countries, such as China, have declared Bitcoin to be a “virtual commodity,” which means that banks are no longer able to exchange Bitcoin on the market. This functions as an implicit ban on digital currencies, an avenue that other countries have taken as well (Thailand and Indonesia). In some countries, including Vietnam and Pakistan, Bitcoin and other cryptocurrencies are banned outright (Global Legal Research Directorate 2018).

Implementation Issues: Once regulators decide along which lines to regulate digital currencies, the implementation has also proven to be difficult due to their decentralized nature. It would be difficult to regulate the actions of every user on the network (or even identify them, given the high level of anonymity). Most attempts at regulating digital currencies have therefore focused on regulating the firms that deal with them, such as businesses and wallet providers. One prominent example is the BitLicense, which businesses dealing with cryptocurrency are required to obtain to operate in New York City (UK Government Chief Scientific Adviser 2016).

The UK’s Office of Government for Science (UK Government Chief Scientific Adviser 2016) emphasizes the need to consider both legal and technical code in the regulation of digital currencies and blockchain technologies. Currently, digital currencies function exclusively through technical code, which determines what is allowed within the network. The rules that govern activity in the network are therefore hardwired into the system and

⁴ The G20 summit that met in Argentina on March 19–20, 2018, highlighted the necessity for a holistic and global approach to regulating cryptocurrencies but concluded that additional data are needed before any recommendations can be put forth.

cannot be broken. This enhances compliance but potentially increases the need for enforcement. The report recommends regulating DLT technologies with a mix of technical and legal code and incorporating public feedback and discussion in the maintenance and creation of the technical code that governs digital currencies.

Market Concerns: From market viewpoints, market participants have been concerned about regulatory uncertainties and the potential interactions of the various regulatory guidance and rules. Middlebrook and Hughes (2014) identify that, in the United States, cryptocurrencies are being regulated by both the state and the federal regulatory authorities as well as by associated foreign governments. Auer and Claessens (2018) find that news related to general bans on cryptocurrencies or their treatment under securities law have adverse effects on cryptocurrency value, whereas news of the establishment of legal frameworks tailored to cryptocurrencies coincide with strong market gains.

Stablecoins: A stablecoin is a type of digital currency backed by a basket of major currency assets. Facebook has attempted to launch a stablecoin through the Libra Association.⁵ The association published a white paper for establishing a dialogue on a stablecoin, Libra, in June 2019. The dialogue with central bankers, regulators, elected officials, and various other stakeholders aimed to determine the best way to combine blockchain technology (Libra Blockchain) with accepted regulatory frameworks. The objective is for the Libra payment system to integrate smoothly with local monetary and macroprudential policies and complement existing currencies by enabling new functionality, drastically reducing costs, and fostering financial inclusion. A large potential userbase could imply fast adoption once the stablecoin is launched. However, the Libra plan generated strong reactions by the international regulatory community, reflecting potentially wide implications for monetary sovereignty, financial stability, and financial integrity.

On April 16, 2020, the Libra Association published a revised version of its 2019 white paper. The revisions attempt to address major concerns raised

⁵ The Libra Association is an independent membership organization headquartered in Geneva, Switzerland. The Association's mission is to provide a governance framework and coordinate decision-making for the Libra network and Libra Reserve, oversee the operation and evolution of the Libra global payment system, facilitate the provision of financial infrastructure and services in a safe and compliant framework through the Libra Blockchain, and support financial inclusion. While Facebook teams played a key role in the creation of the Association and the Libra Blockchain, they have no special rights within the Association. On October 14, 2019, the initial Association Members signed on the Association Council, which is made up of one representative per Member organization (membership of the Association consists of geographically distributed and diverse businesses and nonprofit organizations). Several of the initial 28 members left the association in late 2019/early 2020 in response to serious concerns from financial regulators.

by the international regulatory community during the Association's extensive consultations for an open and trustworthy Libra payment system. The updated white paper makes four key changes below, and the Libra Association was later transformed into Diem Association on December 1, 2020 to offer single-currency stablecoins "Diem Coins" for realtime cross-border payment.

- Offering single-currency stablecoins in addition to the multicurrency coin. Specifically, the design of the new Diem payment system is envisioned to start with single-currency (USD, EUR, GBP, and SGD) coins in addition to the basket currency version. The Association stated that it would welcome the oversight and control over the basket composition (both currencies included and their respective weights) by a group of regulators and central banks or an international organization.
- Enhancing the safety of the Diem payment system with a robust compliance framework. The Association plans to develop strong Anti-Money Laundering and Combating the Financing of Terrorism (AML/CFT) standards, and the establishment of a Financial Intelligence Function to help support and uphold operating standards for network participants.
- Forgoing the future transition to a permissionless system, while maintaining its key economic properties. Regulators have been concerned about the risk in such an open system of unknown participants taking control of the system and removing key compliance provisions. Instead, network participants (service providers) will be selected through a competitive process based on a set of published criteria.
- Building strong protections into the design of the Diem Reserve. The Diem coins will be fully backed by reserves, of which 80% will be invested in liquid short-term government securities, and the remainder will be held in cash. In addition, with input from regulators, the Association will develop a regulatory capital framework to ensure it maintains an appropriately sized, loss-absorbing capital buffer.

3 Central Bank Digital Currency

Central Bank Digital Currency (CBDC) is a digital form of currency issued by central banks that acts as legal tender for the general public. It is meant to serve all the same functions as money but without any physical form. The actual characteristics of a CBDC—for example, how it is held and how transactions are processed—will greatly depend on its design. The analysis here is

by no means exhaustive of the complex web of potential effects of a CBDC. Rather, we highlight several key questions and issues that have emerged in the literature as important considerations when evaluating whether a central bank may wish to issue its own digital currency. For a more thorough analysis, see Mancini-Griffoli et al. (2018) and Bank for International Settlements (2018). For a literature survey on CBDC and cryptocurrencies, see Allen, Gu, and Jagtiani (2021). We highlight several important country experiences with CBDCs in this section.

CBDCs can be classified into two broad types with distinctive characteristics: those based on accounts or those based on tokens. The former would involve a transfer of a claim from one account to another, while the latter would involve a transfer of a token from a digital wallet to another (Mancini-Griffoli et al. 2018). With an account-based CBDC, users would hold accounts directly with a central bank, and payments would involve a user transferring funds from one central bank account to another, similar to today's transactions between commercial bank depositors. A token-based system, on the other hand, would function very similar to cash, in that users could pay other users from their digital wallets directly. These transactions would still need to be verified, however, to ensure that the digital money is not counterfeit or being used more than once. In the account-based case, the process would obviously be highly centralized because all accounts in digital currency would be held at a central bank.

Unlike an account-based system, a token-based system could be either decentralized, using distributed ledger technology to verify and record transactions, or centralized, with the central bank maintaining a central ledger to record and verify transactions (Mancini-Griffoli et al. 2018). The choice of system will have significant implications for the potential anonymity of transactions using digital currency. The exact details of the CBDC design and system have important implications for its adoption and its effects on the larger economy. Although there are a vast number of design choices to be made, several have emerged in the literature as being important in determining the demand for and effects of CBDCs (discussed in greater detail next); see Bank for International Settlements (2018) and Mancini-Griffoli et al. (2018). Next, we discuss some pertinent issues relating to the decisions about and the operations of CBDCs.

3.1 Should Central Bank Issue Digital Currency?

The rapid growth of private cryptocurrencies has ignited discussions around whether central banks should develop their own digital currencies. There

have been concerns around the rapid growth in the adoption of private digital currencies, including their potential destabilizing effects; see Bordo and Levin (2017) and Raskin and Yermack (2016). However, research so far has produced mixed results. Ali et al. (2014) argue that this warning may be somewhat overblown. CBDCs could also potentially improve the conduct of monetary policy, depending on the design, through reducing the effective lower bound on interest rates. However, a recent report by the Central Bank of Canada does not find this to be a compelling reason because there are other less disruptive avenues to reduce the effective lower bound, such as eliminating large currency notes; see Engert and Fung (2017). Overall, CBDC proponents cite increases in the efficiency of payment systems and improved control over monetary policy as important reasons to develop CBDCs, while critics raise concerns about the safety of cryptocurrencies for the financial system.

The most important concerns seem to be related to a potential system-wide run on the banking system toward CBDCs when there is a doubt at private banks, thus potentially increasing instability in the financial system. The creation of CBDCs would cause commercial banks to lose deposits and payment fee revenue. And CBDCs would make it very easy for consumers to withdraw money in the event of a panic.

Some also argue that central banks should create their own digital currencies simply to halt the potentially disruptive rise of private digital currencies such as Bitcoin. A few countries, such as Ecuador, have already run pilot schemes issuing CBDCs to the general public, but the potential costs and effects of issuing CBDCs are still largely unknown. Lacking total agreement around the impact of CBDCs versus private digital currencies, we note that, at the minimum, CBDCs could effectively mitigate the risk of private digital currencies that are growing too large to manage.

CBDC Benefits to Consumers: CBDCs could potentially benefit consumers through increased financial inclusion and the enhanced payment system efficiency. CBDCs represent an efficient way to access, store, and save money, and thus, they could widen access to finance for traditionally underserved and unbanked populations. This is particularly true for developing countries or countries with unstable currencies. Digital currencies also have the potential to increase efficiency in the payments system (by reducing the costs for those transactions that are already undertaken and by enabling transactions that are currently foregone).⁶

⁶ Fung and Halaburda (2016) argue that the main improvement in efficiency would come from reducing costs of current transactions because the size of foregone transactions is relatively small.

Lower Transaction Costs: At present, digital currencies (including private digital currencies) tend to offer lower transaction fees than traditional credit card payments. However, this trend may not be sustainable for private digital currencies. A study by the Bank of England (Ali et al. 2014) argues that private digital currency schemes are typically characterized by increasing marginal cost of producing currency, in part because of the decentralized nature of production. The large and ever-increasing computing power required to create new currency would be difficult for one miner to acquire, except in the case of a severely concentrated market (monopoly or oligopoly), so that marginal costs tend to increase over time. This would tend to erode the efficiencies provided by private digital currencies and giving CBDCs a clear advantage over private digital currencies because CBDCs' centralized production could keep marginal costs low. Barrdear and Kumhof (2016) estimate that a central bank issuance of digital currency of 30% of gross domestic product (GDP) would permanently raise GDP by 3%, in part due to improved efficiency of monetary transactions.

CBDC Design—Anonymity: Anonymity is an important design decision for CBDCs because it represents a benefit for consumers (privacy protection) but a potentially major cost for the system in terms of reduced financial integrity (illegal transactions made easier). In principle, a decentralized token-based CBDC could provide the same level of anonymity as cash does currently, allowing parties to transact without any record or verification. However, this system could lead to significant safety and financial integrity concerns because it would lend itself very easily to criminal activity. This issue could be mitigated through careful planning and appropriate policies, but it nonetheless represents a major challenge. On the other end of the spectrum, an account-based system with accounts held directly at a central bank would allow for much less anonymity of transactions, since a central authority would have detailed records of every transaction that takes place and presumably the identity of the transacting parties. The level of anonymity is therefore a key aspect to take into consideration in the design of CBDCs.

CBDC Design—Transaction Limits: Imposing strict transaction limits on digital wallets would help to mitigate many of the safety and integrity concerns associated with more anonymous CBDCs because it would tend to reduce illegal or fraudulent transactions. However, this would harm the scalability of CBDCs, making it harder for larger businesses or banks, for example, to fully realize the potential benefits.

CBDC Design—Interest Bearing: The potential of CBDCs to allow for interest payments has been widely discussed in the literature, since it is a

major departure from the current system. Both token-based and account-based CBDCs could theoretically pay interest to holders of digital currency; Sweden has demonstrated with its e-krona. However, the logistics of doing so with an account-based system are much simpler, so most literature has assumed that only account-based CBDCs would pay interest. Having a very liquid currency that pays interest and is available to the general public is a significant departure from what exists currently. So whether CBDCs pay interest is a very important design decision and could have important effects on the financial sector overall.

CBDCs' Impact on Central Banks: CBDCs would reduce costs to central banks associated with issuing printing and managing cash, which has been estimated to be 0.5% of GDP in the Euro area; see Hasan et al. (2013). There are, however, big set-up costs to central banks associated with CBDCs, which would require substantial investment in new infrastructure to create, distribute, and maintain the system (Wadsworth 2018). A substantial infrastructure may be needed to mitigate the risk of operational failure and cyberattacks. Also, issuing a digital currency might require central banks to monitor the currency for fraud and terrorist activity, creating monitoring and compliance costs. Digital currencies could easily cross borders giving central banks more to worry about. Also, digital currencies rely on electricity, leaving them vulnerable to electrical outages.⁷ There also may be potential social costs associated with the risk of financial disintermediation; see Bordo and Levin (2017) and Mancini-Griffoli et al. (2018).

Country-Specific Reactions: Overall, the reactions of central banks around the world to the idea of issuing CBDCs have varied. Many countries have undertaken significant research on the potential of issuing digital currency in their specific jurisdictions, including Canada, Norway, Sweden, England, and China, with varying conclusions. Some countries have rejected the idea because of the minimal benefits and potentially high risks. The Reserve Bank of Australia, for example, indicated as early as 2017 that it is not actively pursuing the idea of issuing CBDCs because of a predicted lack of demand and potential financial stability concerns (Richards 2018). New Zealand and the European Central Bank have made similar pronouncements.

Other countries are still exploring the idea with active research agendas focused on CBDCs. The Central Bank of Canada, for example, has published a number of reports examining the potential benefits and costs of issuing

⁷ Building on electricity, digital currencies require a high amount of energy to validate transactions, and payments initiated on the blockchain typically have a delay of 10 minutes, meaning that buyers and sellers would have to wait 10 minutes after a transaction to exchange goods. Moreover, blockchain is not scalable to a high volume of payments because of the computing power and time delay required to validate transactions.

CBDCs and the relevant questions to consider (Fung and Halaburda 2016; Engert and Fung 2017). Their preliminary work emphasizes improved efficiency of payment systems as a potentially important reason to establish CBDCs, but they note that the costs of setting up the system need to be better understood as well as the potential effects on financial intermediation. Norway has produced similar research (Norges Bank 2018).

Sweden has seriously considered issuing its own digital currency: the e-krona. The currency would complement cash and is intended to be used mostly for small retail payments. The Swedish central bank, the Riksbank, is currently undertaking a three-phase project to evaluate the feasibility and effects of CBDCs; the project is due to be completed in late 2019, at which point the Riksbank will decide whether to move forward with CBDCs.

There is somewhat more enthusiasm for CBDCs among emerging economies with less developed payment systems that typically cite increasing financial inclusion as a main driver for exploring digital currency (Mancini-Griffoli et al. 2018). China is perhaps the major country closest to releasing a publicly available digital currency. The country recently launched a digital currency trial in four urban centers of the country with a combined population of almost 40 million. The government will be using digital currency for certain services in these areas, such as providing transit subsidies (Chaudhury 2020). Furthermore, China has also established a Digital Currency Research Lab within the central bank. Over 40 patent applications have been filed with China's State Intellectual Property Office related to technological innovations for issuing digital currency (Zhao 2018).

Ecuador is the most notable example where a central bank issued a digital currency that was widely available to the public. Starting in 2014, Ecuadorians could open accounts with the central bank in digital currency and transfer these funds to other accounts using a mobile app. The reason behind the creation of the digital currency was to increase financial inclusion and reduce the cost of printing currency (Rosenfeld 2015). The digital currency was subsequently decommissioned after three years because of its failure to attract a significant number of users and the low volume of transactions. This underlines the fact that the potential demand for CBDCs is an important first consideration.

4 Initial Coin Offering (ICO)

ICOs are mechanisms to raise funds by selling coins or tokens, using blockchain technology, to support a product launch or a new virtual currency.

In this way, ICOs are a conjunction of crowdfunding and blockchain. The tokens purchased in an ICO give the participant certain rights, most frequently the right to use the platform services that are being developed through the project and less often profit or ownership rights. The coins can also be exchanged for other cryptocurrencies (and even potentially fiat currencies) on secondary markets. They operate similarly to initial public offerings (IPOs) but typically skirt the usual regulations and restrictions on IPOs. While start-ups have traditionally relied on venture capital to raise funds and grow, ICOs present a more decentralized and democratic alternative. ICOs have offered a significant opportunity for gaining access to capital for small businesses and fintech start-ups.

There have been questions around the legality and trading of ICOs; see the SEC Chairman's remark (Clayton 2017). The use of ICOs, beginning as a relatively unknown fundraising method used only in blockchain circles, has increased dramatically since 2013.⁸ According to PwC, ICOs of cryptocurrencies have raised over US\$31 billion through late 2019 with more than 2000 projects funded (PwC 2020). The United States, Hong Kong, Singapore, and the United Kingdom are the leading countries in terms of the number of completed ICOs.

Fraud and consumer protection concerns have notably led some countries, including China, South Korea, and Pakistan, to ban ICOs altogether (Global Legal Research Directorate 2018). In other countries, ICOs have been issued without consideration of existing regulation (e.g., for securities or crowdfunding). In what circumstances could ICOs be subject to regulation? Is the existing regulatory framework (e.g., for securities or crowdfunding) appropriate for ICOs? In most instances, ICOs raise particular challenges for regulators because they do not usually play by the traditional rules, and the legal classification of a coin or token may depend on the specific circumstances. The following issues are pertinent to the proper functioning of ICOs.

Information Asymmetry Issue: There are several notable issues with ICOs that have caught the attention of regulators. Perhaps most prominently, ICOs are widely seen as profiting from significant information asymmetry between investors and the organizers. Based on a taxonomy of ICOs, for example, it was found that over 20% of the offerings included no information about the initiators or backers of the project, and 25% include no information about the project's financial plans, such as how the funds will be used and in what stages (Zetzsche et al. 2018). The fact that investors continue to fund projects

⁸ For more in-depth evaluating framework for ICOs, see Li and Mann (2018).

with very little information about the underlying investment highlights the potential bubble-like characteristics of the ICO boom, a sort of irrational exuberance. And in the worst cases, ICOs may be merely a front for fraud.

Similarly, Momtaz (2020) finds significant evidence of moral hazard in ICO offerings. The author shows that token issuers systematically exaggerate information in their white papers to impress investors, and this tactic tends to work, as issuers who exaggerate raise more funds in less time compared to those who are more accurate in their white paper details. In the longer term, however, this strategy does backfire as investors learn of the bias through trading with other investors. This highlights a potential need for regulation to ensure investors are properly informed on the potential risks and rewards of an individual token offering.

Security Issues: The paucity of relevant information provided for many ICOs means that investors often have very weak legal protection. Without knowledge of who is backing or issuing an ICO, it is difficult to pursue any sort of legal action against them, which significantly limits investors' possibilities of reclaiming lost or stolen funds (Zetsche et al. 2018). ICOs also may present significant security issues, a fact highlighted by the case of the DAO project, a high-profile ICO that raised \$150 million, \$60 million of which was eventually stolen because of a problem in the code (Del Castillo 2016).

Given the currently thin level of regulatory oversight, ICOs and digital currencies pose significant risks for investors owing to the high potential for fraud and dramatic price swings associated with an unregulated market. Although the impacts continue to be small and localized, the lack of consumer protections related to digital currencies is also a cause for concern. The potential for digital currencies to cloak illegal activities is problematic for the integrity of the financial system. Research has highlighted that, for digital currencies to survive, they must be able to operate in a more regulated framework (De Filippi 2014; Middlebrook and Hughes 2014).

Regulation of ICOs: ICO regulations vary widely across countries, ranging from quite permissive (Switzerland) to outright bans (China). Switzerland is undoubtedly the country that has the most developed regulatory framework. FINMA, Switzerland's financial watchdog, has approved blockchain innovation and dedicated fintech licenses and has instituted ICO regulations that depend on how the ICO is categorized.⁹ FINMA has also shown that they are willing to strictly apply traditional market rules and regulations to ICOs,

⁹ In the Netherlands, regulation depends on whether a token offered is considered a security or a unit of collective investment. The token offered is categorized as a debt security, equity security, managed investment product, or derivative.

as evidenced by actions taken against firms operating ICOs without a license (Wilson 2019).

Even within the same country, ICO regulation and the institutions responsible vary depending on how the ICO is categorized (e.g., as a debt security, equity security). And, as mentioned earlier, some countries have banned ICOs altogether. Bringing cryptocurrencies into the regulated space will increase accountability and hence build trust in them as an investment vehicle or means of payment, which is critical to their survival. The manner in which cryptocurrencies will be regulated is still an open question and will also depend on how they are used in specific instances.

Given the decentralized nature of ICOs, many have noted that the path forward for the regulation of ICOs should involve integrating legal and regulatory principles directly into the code that governs the platforms operations, as this would likely be the best mechanism for ensuring safeguards are in place (Robinson II 2018; UK Government Chief Scientific Adviser 2016). This would require governments to work directly with platform developers, such as Bitcoin and Ethereum, to hardwire the system with investor protections.

Transparency: The vast majority of countries have yet to develop a coherent framework for regulating ICOs. Most countries have issued warnings about the risks of investing in ICOs, highlighting the unregulated and risky nature of the market. In the United States, SEC Chair Jay Clayton (2017) has outlined important considerations for market participants when considering investing in ICOs. Clayton stresses the need for investors to ask questions about the investment opportunity, such as where the money is going, what it is funding exactly, what the rights of the token are, and with whom they are contracting to fully understand the risks and potential for profit. Although these seem like basic questions, it is stunning to see the number of ICO white papers that do not include this basic information. Based on a taxonomy of ICOs, for example, it was found that over 20% of the offerings included no information about the initiators or backers of the project (Zetzsche et al. 2018).

ICOs Under Security Regulations: Although the regulatory environment remains murky, the SEC has applied securities laws to some ICOs that meet the definition of securities. For example, after an investigation into the DAO (decentralized autonomous organization) ICO described earlier, the SEC ruled that the tokens distributed as part of this ICO represent securities, and as such, were subject to the U.S. Securities Exchange Act (Securities and Exchange Commission 2017). No enforcement action was taken against the DAO organizers; nevertheless, the ruling sent a signal to the investors that

the relevant authorities will enforce existing regulations for ICOs when applicable. It seems reasonable to assume that digital currencies provided through ICOs would be regulated as securities, and this seems to be the general trend, although many governments have indicated they will regulate ICOs on a case-by-case basis, owing to the wide variety of forms they take. Regulators in many Asian countries, however, have drafted specific regulations for ICOs (Wockener et al. 2017).

The above ruling makes it clear that ICOs will be increasingly subject to U.S. securities laws, and therefore the regulatory framework for the future will most likely be one of fitting ICOs into existing regulations, with a focus on stopping fraudulent ICOs. In terms of regulation for the future, ICO originators and intermediaries should be required to disclose a minimum of information regarding the project so that investors can make informed decisions. The required information should include the ICO backer's contact information, a description of how investor money will be spent and how it will be stored (by intermediaries; and if so, which ones) as well as an explicit description of the rights conferred with the token purchase.

5 Investing in Cryptocurrencies

Cryptocurrencies have also been used as investment vehicles, thus being referred to as crypto assets—with distinct risk-return characteristics. Although detailed investment information is only available for the two main cryptocurrencies; namely, Bitcoin and Ethereum, investments through ICOs have been registered for more than 222 cryptocurrencies (out of over 1300), with a total market capitalization of approximately US\$250 billion over 2013–2017 (where the main share is in Bitcoin) and a market capitalization of over \$1 trillion for Bitcoin alone as of April 2021. The purpose of raised capital through ICOs ranges from infrastructure financing to peer-to-peer (P2P) lending and crowdfunding. A report from Bloomberg (Qui et al. 2018) finds that cryptocurrencies have weak relationships with established asset classes. Still, cryptocurrencies have strong correlations with other cryptocurrencies, so their prices tend to move together.

Similarly, Hu et al. (2018) investigate a common variation for secondary market returns on 222 digital coins. They find that the secondary market returns of all other cryptocurrency returns co-move with Bitcoin returns but are not well correlated with traditional assets returns. Similarly, Liu and Tsyvinski (2018) find that cryptocurrency returns are largely predicted by factors specific to cryptocurrency markets, having little to no exposure to

most stock market and macroeconomic factors. However, one should be cautious about thinking of cryptocurrency as a diversification tool for their portfolio. The current regulatory uncertainty relating to cryptocurrencies has remained a major concern among investors around the globe. The returns yielded by digital assets can be high in some cases, but it is critical that investors understand the risks.

Unlike traditional assets that prompt a fundamental analysis of an industry or company (business plan, financials, and leadership) while considering many other factors, cryptocurrency investments are speculative and the returns are extremely volatile and largely unregulated. Clearer and more stringent regulation of ICOs and cryptocurrencies in general will reduce the risk of investor fraud and make cryptocurrencies more attractive investment options for general investors. This should spur more growth in the sector and bring added benefits to investors.

Who Could Be Investors? One would expect institutional investors such as hedge funds, traditional asset managers, and venture capital to be the primary investors in this market. Anecdotal evidence suggests that the majority of cryptocurrency investors are small-scale investors with a rather short-term investment horizon. Identification, measurement, and availability of instruments to manage cryptocurrency risks, especially relating to price volatility and safety of assets, remain a major challenge for investors. There are limited hedging tools for cryptocurrencies, for example, options for Bitcoin are relatively expensive because of its high volatility. In addition, there have been concerns around performance measurement that lacks standardization around benchmarks and token valuation.

Asset Bubble: Many investors have profited significantly from Bitcoin and other cryptocurrencies. Indeed, an investor who bought Bitcoin in December 2015 and sold it two years later in December 2017 would have seen a return over 4300%, a simply eye-popping figure (Coinbase 2018). However, this statistic hides significant volatility, and this dramatic increase can only be characterized as a bubble, a bubble that eventually burst. Bitcoin's price has fallen 82% in 2018, and 38% in November alone, to a price of around \$3200 as of December 2018 (Coinbase 2018), a fact that underlines the riskiness of cryptocurrency investment. As mentioned earlier, the value of Bitcoin recently perked at over \$63,000 in April 2021. Stories abound about ordinary consumers who lost significant sums of money in the volatile cryptocurrency markets (Popper and Lee 2018).

Speculative Vehicle: Most recent research concludes that Bitcoin and other cryptocurrencies bear many of the characteristics of a speculative vehicle, rather than an investment vehicle. Baek and Elbeck (2015) find that the

Bitcoin market is highly risky and speculative relative to the stock market. Similarly, Klein et al. (2018) predict that Bitcoin will continue to experience large price drops as investors continue to profit at the peak of price movements. Baur et al. (2018) come to a similar conclusion; however, they note that Bitcoin's particular return properties, which are very different from traditional asset classes, make Bitcoin a good diversification investment. They also posit that, as use of Bitcoin grows, its price volatility will drop, and it will become a more balanced investment vehicle.

Demand-Driven Aspect of Cryptocurrencies: As we saw earlier, the price of Bitcoin and other cryptocurrencies is a function of people's beliefs about future supply and demand for the currency. Since supply is largely fixed in the long term for most cryptocurrencies, the price will be reflective mainly of expectations of demand in the future. These expectations, in turn, should largely be formed by beliefs about the future use of cryptocurrencies as a medium of exchange that will continue to remain in demand into the far future. Research has shown this is somewhat unlikely to be the case (Ali et al. 2014), so investors should not be surprised to see the price of cryptocurrencies gradually decline over time as users and investors realize their limited value. Barring improvements in technology that make digital currencies more viable alternatives to traditional currencies, cryptocurrencies do not appear to be good investments in the long term.

ICOs, on the other hand, represent an important source of funding (when designed properly) for new and innovative firms that may not have access to capital elsewhere. ICOs have grown dramatically in recent years, having raised an estimated \$7 billion since 2017 (Catalini and Gans 2018). To date, ICOs have mainly been used to raise funding for blockchain-based technology applications, and this trend is likely to continue. The majority of ICOs still provide mainly "utility" tokens, which grant access to platform services and other products funded through the ICO, and therefore differ significantly from traditional securities or investments. Indeed, almost 70% of ICOs offer access to platform services as a major right, while only 26 and 21% provide rights to profits and fiat currency, respectively (Adhami et al. 2018).

Caution: Although ICOs have served as a good source of fund for new fintech start-ups, experience has shown the need for extreme caution when investing in these ventures. Indeed, the U. S. SEC has issued strong warnings about investing in ICOs, highlighting the increased potential for fraud because of the lack of regulations. It is also important to note that ICO markets are international, which may prevent the relevant authorities from

pursuing corrective or pecuniary action. This further heightens the risks ICOs posed to investors (Clayton 2017).

Overall, the case for cryptocurrencies and ICOs as investment vehicles is mixed. On the one hand, significant volatility and a lack of regulation make them risky investments, with a high potential for losses, owing to market swings, poor business plans, and potentially outright fraud. On the other hand, ICOs may help channel funds into highly productive and innovative ventures that otherwise might not be funded. Whatever the case may be, it is clear that increased and improved regulation of these markets is needed to ensure investors are protected and that the innovations are encouraged.

6 Concluding Remarks

Cryptocurrencies, as an alternative to national currencies as a medium of exchange and store of value, have the potential to reduce cost and increase efficiency in the payments system. Their ability to do so will largely depend on whether cryptocurrencies can continue to promise low transaction fees and ease of use into the future, which research has shown may not be the case (Ali et al. 2014). But with improvements in technology, digital currencies may be able to maintain their low transaction fees and thus provide significant benefits to consumers and businesses. Few predict that digital currencies will completely replace traditional ones; rather, they are destined to coexist with traditional legal tenders, adding value in currently underdeveloped areas such as in international money transfers.

There have been concerns that digital currencies will lead to a loss of control over monetary policy and significant disruption to the financial sector, which appear to be largely overblown. Research has generally concluded that, at their current level, digital currencies have a miniscule impact on monetary policy; usage would have to increase dramatically to pose any real threat, a scenario that is highly unlikely by most accounts. The conclusion is similar when considering financial stability.

Although the use of cryptocurrencies is not seen to pose any monetary policy or financial stability concerns, national and international monetary and financial authorities need to remain vigilant about the resultant risks from their evolving uses. Further, they should be ready to introduce regulation at the national and international level to curb hazards stemming from possible payment and clearing system disruptions and from transactions of questionable integrity. At the same time, authorities should usefully take advantage of any new DLT that could improve the efficiency of the existing payment and

clearing systems. Whatever the effect of digital currencies, it is clear that the technology underpinning their use (such as blockchain and other DLT technologies) will have profound effects in finance and payments, with dramatic potential to improve efficiency and transparency in these areas.¹⁰

References

- Allen, Franklin, Xian Gu, and Julapa Jagtiani. 2021. A Survey of Fintech Research and Policy Discussion. *Review of Corporate Finance*, 1: 1–81.
- Adhami, Saman, Giancarlo Giudici, and Stefano Martinazzi. 2018. Why do businesses go crypto? An empirical analysis of initial coin offerings. *Journal of Economics and Business* 100: 64–75.
- Ali, Robleh, John Barrdear, Roger Clews, and James Southgate. 2014. The economics of digital currencies. *Bank of England Quarterly Bulletin* 2014: Q3.
- Auer, Raphael, and Stijn Claessens. 2018. Regulating cryptocurrencies: Assessing market reactions. *BIS Quarterly Review*. https://www.bis.org/publ/qtrpdf/r_qt1809f.pdf.
- Baek, C., and M. Elbeck. 2015. Bitcoins as an investment or speculative vehicle? A first look. *Applied Economics Letters* 22 (1): 30–34.
- Ball, James, Charles Arthur, and Adam Gabbatt. 2013. FBI claims largest Bitcoin seizure after arrest of alleged silk road founder. *The Guardian*, October 2.
- Bank for International Settlements. 2018. Central Bank Digital Currencies. Committee on Payments and Market Infrastructures, Markets Committee. <https://www.bis.org/cpmi/publ/d174.pdf>.
- Barrdear, John, and Michael Kumhof. 2016. The macroeconomics of central bank issued digital currencies. Bank of England Staff Working Paper 605. <https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2016/the-macroeconomics-of-central-bank-issued-digital-currencies.pdf?la=en&hash=341B602838707E5D6FC26884588C912A721B1DC1>.
- Baur, Dirk G., KiHoon Hong, and Adrian D. Lee. 2018. Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions, and Money* 54: 177–189.
- Blockchain.com. 2018. Confirmed transactions per day. <https://www.blockchain.com/charts/n-transactions>. Accessed in December 2018.
- Bordo, Michael D., and Andrew T. Levin. 2017. Central bank digital currency and the future of monetary policy. NBER Working Paper Series 23711. Cambridge, MA: NBER.

¹⁰ The DLT also holds considerable promise for governments—with the potential to help governments to collect taxes, deliver benefits, issue passports, record land registries, assure the supply chain of goods, and generally ensure the integrity of government records and services; see UK Government Chief Scientific Adviser (2016: 6). This is illustrated by the case of E-Estonia, an effort by the Government of Estonia to digitize the vast majority of its services, which has seen impressive results (Heller 2017).

- Catalini, Christian, and Joshua S. Gans. 2018. Initial coin offerings and the value of crypto tokens. NBER Working Paper Series 24418. Cambridge, MA: NBER.
- Chaudhury, Dipanjan Roy. 2020. China flirts with digital currency in four urban centers amid Covid. *The Economic Times*, May 2020. <https://economictimes.indiatimes.com/news/international/business/china-flirts-with-digital-currency-in-four-urban-centres-amid-covid/articleshow/75504372.cms>.
- Clayton, Jay. 2017. Statement on cryptocurrencies and initial coin offerings, December 11. https://conferences.law.stanford.edu/vcs/wp-content/uploads/sites/11/2017/11/SEC.gov_-_Statement-on-Cryptocurrencies-and-Initial-Coin-Offerings.pdf.
- Coinbase. 2018. Cryptocurrency price charts. <https://www.coinbase.com/charts>. Accessed in December 2018.
- CoinDesk. 2014. State of Bitcoin 2014. *Wall Street Journal*, February. <http://online.wsj.com/public/resources/documents/CoinDeskStateofBitcoin.pdf>.
- CoinLore. 2020. Cryptocurrency list. https://www.coinlore.com/all_coins. Accessed in May 2020.
- De Filippi, Primavera. 2014. Bitcoin: A regulatory nightmare to a libertarian dream. *Internet Policy Review* 3 (2). <https://policyreview.info/node/286/pdf>.
- Del Castillo, Michael. 2016. Coindesk.com, June 18. <https://www.coindesk.com/the-dao-an-analysis-of-the-fallout>.
- Engert, Walter, and Ben S.C. Fung. 2017. Central bank digital currency: Motivation and implications. Bank of Canada Staff Discussion Paper. <https://www.bankofcanada.ca/wp-content/uploads/2017/11/sdp2017-16.pdf>.
- Federal Reserve System. 2017. Federal reserve payments study 2016. Federal Reserve System. <https://www.federalreserve.gov/newsevents/press/other/2016-payments-study-20161222.pdf>.
- Forbes Finance Council. 2017. Eight ways digital currencies will change the financial landscape. *Forbes*, August 25. <https://www.forbes.com/sites/forbesfinancecouncil/2017/08/25/eight-ways-digital-currencies-will-change-the-financial-landscape/#37a04e382304>.
- Fung, Ben S.C., and Hanna Halaburda. 2016. Central bank digital currencies: A framework for assessing why and how. Bank of Canada Staff Discussion Paper.
- Global Legal Research Directorate. 2018. Regulation of cryptocurrency around the world. The Law Library of Congress, Global Legal Research Center. <https://permanent.access.gpo.gov/gpo94288/cryptocurrency-world-survey.pdf>.
- Hale, Galina, Arvind Krishnamurthy, Marianna Kudlyak, and Patrick Shultz. 2018. How futures trading changed bitcoin prices. Federal Reserve Bank of San Francisco Economic Letter.
- Hasan Iftexhar, Tania De Renzis, and Heiko Schmiedel. 2013. Retail payments and the real economy. European Central Bank Working Paper 1572. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1572.pdf>.
- Heller, Nathan. 2017. Estonia, the digital republic. *The New Yorker*, December 18 and 25, 2017. <https://www.newyorker.com/magazine/2017/12/18/estonia-the-digital-republic>.

- Hu, Albert, Christine Parlour, and Uday Rajan. 2018. Cryptocurrencies: Stylized facts on a new investible instrument. University of California, Berkeley Working Paper (June 5, 2018), available on SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3182113.
- Irrera, Anna, and Brenna Hughes Neghaiwi. 2018. Switzerland seeks to regain cryptocurrency crown. *Reuters*, July. <https://www.reuters.com/article/us-cryptocurrencies-banking-switzerland/switzerland-seeks-to-regain-cryptocurrency-crown-idUSKBN1K91AY>.
- Jagtiani, Julapa, Michael G. Papaioannou, and George Tsetsekos. 2019. Cryptocurrencies in the global economy. *The Journal of Investing, Cryptocurrency* 28 (3), April 7–12. <https://joi.ijournals.com/content/28/3/7> or <https://joi.ijournals.com/content/28/3>.
- Kharpal, Arjun. 2018. Cryptocurrencies: Regulating the new economy. CNBC, August. <https://www.cnbc.com/2018/08/09/cryptocurrencies-regulating-the-new-economy.html>.
- Klein, Tony, Hien Pham Thu, and Thomas Walther. 2018. Bitcoin is not the New Gold—A comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis* 59: 105–116.
- Li, Jiasun, and William Mann. 2018. Initial coin offerings and platform building (2018). 2018 WFA, 2019 AFA. <https://ssrn.com/abstract=3088726> or <http://dx.doi.org/10.2139/ssrn.3088726>.
- Liu, Yukum, and Aleh Tsyvinski. 2018. Risks and returns of cryptocurrency. NBER Working Paper Series 24877. Cambridge, MA: NBER.
- Libra Association Members. 2019. An introduction to Libra—White Paper, June. https://libra.org/en-US/wp-content/uploads/sites/23/2019/06/LibraWhitePaper_en_US.pdf.
- Libra Association Members. 2020. Libra—White Paper V2.0, April. <https://libra.org/en-US/white-paper/>.
- Mancini-Griffoli, Tommaso, Maria Soledad Martinez Peria, Itai Agur, Anil Air, John Kiff, Adina Popescu, and Celine Rochon. 2018. Casting light on central bank digital currency. IMF Staff Discussion Note. <https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2018/11/13/Casting-Light-on-Central-Bank-Digital-Currencies-46233>.
- Mbogo, Angeline. 2018. Cryptocurrencies can boost financial inclusion experts agree. *Bitcoinafrica.io*. March 20. <https://bitcoinafrica.io/2018/03/20/cryptocurrencies-can-boost-financial-inclusion/>.
- Middlebrook, Stephen T., and Sarah Jane Hughes. 2014. Regulating cryptocurrencies in the United States: Current issues and future directions. *William Mitchell Law Review* 40 (2). <https://open.mitchellhamline.edu/cgi/viewcontent.cgi?article=1567&context=wmlr>.
- Momtaf, Paul. 2020. Entrepreneurial finance and moral Hazard: Evidence from token offerings. *Journal of Business Venturing*.
- Norges Bank. 2018. Central bank digital currencies. *Norges Bank Papers* 1/2018. <https://static.norges-bank.no/contentassets/166efadb3d73419c8c50f9471be>

- 26402/nbpapers-1-2018-centralbankdigitalcurrencies.pdf?v=05/18/2018121950&ft=.pdf.
- Popper, Nathaniel, and Su-Hyun Lee. 2018. After the Bitcoin boom: Hard lessons for cryptocurrency investors. CNBC, August 20. <https://www.cnbc.com/2018/08/20/after-the-bitcoin-boom-hard-lessons-for-cryptocurrency-investors.html>.
- PwC (Strategy&). 2020. 6th ICO/STO Report: Spring 2020 Edition. https://www.pwc.com/ee/et/publications/pub/Strategy&_ICO_STO_Study_Version_Spring_2020.pdf.
- Quarles, Randal K. 2017. Thoughts on prudent innovation in the payment system. 2017 Financial Stability and Fintech Conference. Federal Reserve Bank of Cleveland. November 30.
- Qui, Yue, Justina Lee, and Adrian Leung. 2018. How crypto stacks up against other investments. Bloomberg. <https://www.bloomberg.com/graphics/2018-bitcoin-investment/>.
- Raskin, Max, and David Yermack. 2016. Digital currencies, decentralized ledgers, and the future of central banking. NBER Working Paper Series 22238. Cambridge, MA: NBER.
- Richards, Tony. 2018. Cryptocurrencies and distributed ledger technology. Speech at the Australian Business Economists Briefing, June 26, Sydney, Australia. <https://www.rba.gov.au/speeches/2018/sp-so-2018-06-26.html#fn9>.
- Robinson, I.I., and A. Randolph. 2018. "The new digital wild west: Regulating the explosion of initial coin offerings. *Tennessee Law Review* 85 (4): 897–960.
- Rosenfeld, Everett. 2015. Ecuador becomes the first country to roll out its own digital cash. MSNBC, February, 2015. <https://www.cnbc.com/2015/02/06/ecuador-becomes-the-first-country-to-roll-out-its-own-digital-durrency.html>.
- Securities and Exchange Commission. 2017. Report of investigation pursuant to Section 21(a) of the Securities Exchange Act of 1934: The DAO. Release 81207, July 25. <https://www.sec.gov/litigation/investreport/34-81207.pdf>.
- Securities and Exchange Commission. 2018. Five things you need to know about ICOs. August 21. <https://www.sec.gov/ICO>.
- Standard and Poor's. 2018. S&P Global Ratings. https://www.standardandpoors.com/en_US/web/guest/home. Accessed in December 2018.
- Suberg, William. 2018. G20 forum shelves deadline for 'very specific recommendations' on crypto. *Cointelegraph*, July 23. https://cointelegraph.com/news/g20-forum-shelves-deadline-for-very-specific-recommendations-on-crypto?utm_source=yahoofinance&utm_medium=referral&utm_campaign=content.
- UK Government Chief Scientific Adviser. 2016. Digital ledger technology: Beyond block chain. Government Office for Science, UK. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf.
- Wadsworth, Amber. 2018. The pros and cons of issuing a central bank digital currency. *Reserve Bank of New Zealand Bulletin* 81 (7). Reserve Bank of New Zealand. <https://www.rbnz.govt.nz/-/media/ReserveBank/Files/Publications/Bulletins/2018/2018jun81-07.pdf>.

- Wilson, Tom. 2019. Swiss financial watchdog finds \$90 million initial coin offering illegal. *Reuters*, March.
- Wockener, Karsten, Carsten Losing, Thilo Diehl, and Annkatrin Kutzbach. 2017. Regulation of Initial Coin Offerings. White & Case: Our Thinking, December 15.
- Zetsche, Dirk A., Ross P. Buckley, and Douglas W. Arner, and Linus Föhr. 2018. The ICO gold rush: It's a scam, it's a bubble, it's a super challenge for regulators. *Harvard International Law Journal* 63 (2): 2019. <https://ssrn.com/abstract=3072298> or <http://dx.doi.org/10.2139/ssrn.3072298>.
- Zhao, Wolfie. 2018. PboC filings reveal big picture for planned digital currency. *Coindesk*, June 26. <https://www.coindesk.com/pboc-filings-reveal-big-picture-for-planned-digital-currency>.

Payments and Trading



Fintech Revolutionizing Financial Services: The Case of Virtual Savings and Credit Supply Technological Platforms in Africa

Njuguna Ndung'u and Alex Oguso

1 Introduction

In an innovative way, Africa is solving its financial exclusion problems and charting its own route for more inclusive financial services and driving economic vibrancy with it. This section explores the innovative developments of financial services through the mobile phone transactions platform, embodied by the launch of M-Pesa in Kenya in 2007. This innovative financial service developed further from the confines of money transfer to become a real time retail electronic payments platform. The developments that followed is what may be termed as “success begets success” that pushed this platform to one of virtual savings and virtual credit supply platform. The mobile phone-based transactions accounts have given way to retail electronics

Njuguna Ndung'u is Professor of Economics and the Executive Director, African Economic Research Consortium (AERC); and the immediate former Governor of the Central Bank of Kenya, 2007–2015. Alex Oguso is an Economist & Researcher at Strategy, Innovations and Risk Management Department, Kenya Revenue Authority. All comments and questions to be addressed to nsnjuguna@ndungu.co.ke.

N. Ndung'u (✉)

African Economic Research Consortium, Nairobi, Kenya

e-mail: nsnjuguna@ndungu.co.ke

A. Oguso

Kenya Revenue Authority, Nairobi, Kenya

payments system, digital financial services and to an overall digital revolution that is spreading to several economies in Africa. The section provides examples of successful virtual products in five African economies.

In the period prior to 2007, before the launch of M-Pesa and the ensuing digital revolution in the financial sector in Kenya, obtaining a short-term loan from the banks had many requirements including time. Gubbins and Totolo (2018) showed that it took several days (the 2016 Kenya household budget survey found that on average—across all loan types—borrowers needed 11 days to secure a short-term loan). If it was a long-term loan that required formal security, other than employer's information, it took months for the collateral to be registered, credit committees to meet and lawyers to perfect the collateral charge. Additionally, on top of the high-interest rate charges and collateral requirements, application for a bank loan involved out-of-pocket costs such as physical travel to the point of service, documentation, and applications (Gubbins and Totolo 2018). Some consumers, especially small businesses, could not meet the collateral requirements to access the banks loans. Moreover, the long procedure and time taken to process loans meant that very short-term urgent business needs for funds to bridge the gap between spending cash and receiving funds (Bennett et al. 2016) were not met. This resulted into lost business opportunities and/or continued cash flow problems. Bennett et al. (2016) pointed out that if banking decisions and transactions cannot be completed quickly, then entrepreneurs cannot immediately react to changes in the business environment and latch on to opportunities swiftly as soon as they spot them. Consequently, this leads to a significant use of informal sources of funding such as moneylenders (shylocks), credit terms from suppliers, *chamas*,¹ friends, and family. These informal financial outlets for entrepreneurs are costly, risky, and not conducive to business development. This indirectly seems to endogenously raise the demand for a new avenue of financial services, and an efficient platform of financial services. The digital financial services platform appears then to have gained grounds due to this inadequate and inefficient financial services regime. In addition, it was pervasive across African economies. In this regard, and thanks to the digital evolution that has been with us in the last 13 years now, an efficient platform of financial services is slowly developing in the African region. This is a digital platform of virtual banking services. It is not dependent on physical branch outlets. It is branchless and spreading fast in the continent. We show some examples on how it is functioning and spreading with examples from five countries in Africa.

¹ This refers to a group or an informal cooperative society through which members pool savings and borrow money for their personal use, as well as for individual or group investments.

2 The Enabling Framework

The chronology of events from 2007 after M-Pesa in Kenya was launched is a major illustration of innovation in the financial system that supports the concept of alternative finance in this section. M-Pesa is an electronic money transfer product that enables users to store value on their mobile phone or mobile account in the form of electronic currency that can be used for multiple purposes including transfers to other users, payments for goods and services, and conversion to and from cash. First, M-Pesa introduced a money transfer platform that was efficient and real time, Second, this money transfer platform: Person to Person (P2P), Person to Government (P2G), and Person to Company (P2C²) transfers slowly turned into a real time retail electronic payments platform. This functionality could be effected at the Agent's counter, for those with no telephones but also using the phone at a place of comfort and convenience. So, there was no need to visit a bank or a point of service/or sale to effect electronic payments. The underlying success is the use of electronic units of money, e-money, as the basis for transactions on the mobile phone. This was supported by an electronic transaction platform in a commercial bank and the telecommunication network provided the transmission backbone.

It took a combination of innovative telecommunication and banking sectors and a sound regulatory environment for both to enable the emergence of new financial products for the poor—including many young users (Ndung'u 2019). The four stages of digitization in the financial services evolution in Kenya were facilitated by a “*test and learn*” approach applied by the Central Bank of Kenya, the Communications Authority and the Competition Authority. Kenya is often cited as an example of success driven primarily by restrained regulators who preferred to set rules *ex post* as services and their providers evolved rather than impose a strict *ex ante* regime that would later prove inappropriate for the market (Ndung'u 2019). The celebrated “*test and learn*” approach allowed the telecommunication companies to collaborate with the commercial banks to innovate and roll out various financial products and revolutionize the financial system in Kenya. Ndung'u (2019) pointed out that the financial sector regulators have since embraced regulatory Sandboxes and regulatory technology (RegTech) to help keep up with the substantial changes in financial services marketplaces, which has enhanced their capacity to monitor, evaluate, and ensure compliance in the sector.

² P2P is person to person money transfer, P2G is person to government transfer and vice versa, P2C is person to company transfer and vice versa.

The digital financial services evolution in Kenya was also supported by a comprehensive and reasonably functional system to provide unique identification to its population, which is extensively used for “know-your-customer (KYC)” purposes. The National Registration Bureau collects biographic and biometric data through its application process and operates the system wide Automated Fingerprint Identification System, which checks against possible duplicate or multiple registrations (World Bank 2016). As pointed out by Ndung'u (2019), the Kenyan government is planning to replace the current “second generation” national ID (Identification) cards with a new ID system including a machine-readable card with a microchip that will contain more comprehensive details of the holder. This new card will make it easier to transmit transaction and other data in real time. The government is in the process of implementing a six-character Unique Personal Identifier that will be linked to an electronic database with the educational records of all individuals from primary school up to university level. Bringing the citizens onto the mobile phone-based digital platform involves registration of the SIM cards. Initially, SIM card registration was not required, but when M-Pesa and similar products entered the market, the Central Bank of Kenya required that participants in the mobile phone-based payments ecosystem should have their SIM card registered. Ndung'u (2019) noted that the Kenya Communications Act was amended to regulate SIM card registrations such that a telecommunications operator is required to obtain a customer's full name, identity card number, date of birth, gender, and physical and postal address before selling a SIM card or providing telecommunication services. The law requires that telecommunication operators ensure that proper physical or electronic records are kept of the subscribers' information. It also allows telecommunications operators to disclose the registration particulars of a subscriber in connection with the investigation of any criminal offence or for the purpose of any criminal proceedings. It was then a matter of time that a virtual savings platform was developed—that is the birth of M-Shwari as the virtual savings product. As we show later, it is a revolution in financial services and a reflection of market development.

The launch of M-Shwari in November 2012 and the developments in virtual savings and credit that ensued have transformed the saving and credit supply process in Kenya's financial market. In these developments we also show that transaction data and savings were used to generate credit scores for short-term virtual credit, thus revolutionizing the credit market in Kenya. Looking at the overall credit landscape, Gubbins and Totolo (2018) showed that virtual credits seem to have taken the position of an all-purpose borrowing tool with a broad appeal. The virtual credits are processed,

disbursed and repaid remotely, without the need for customers to visit bank branches or outlets. Loan eligibility decisions are automated based on a set of rules applied to available data, thus the virtual credits are instant (Gubbins and Totolo 2018). At the beginning of M-Shwari, it was clear that savings data, transactions data, and the tiered KYC (Know Your Customer) would be used to safeguard the credit supply platform in Kenya as well as generate appropriate credit scores to enable a pricing mechanism of short-term credit. This is an innovation not known before. Evidently, the digital financial platform has transformed the intermediation process between the borrowers and lenders. This has offered an opportunity for alternative sources of finance, reviewed in this section.

The rest of the section is organized as follows: the next part of the section looks at the digital evolution that has taken place in Kenya, followed with a discussion on how this digital evolution has pushed frontiers of financial inclusion. The section then presents the virtual savings and credit supply products developing in Africa. Lastly, the section summarizes the developments so far witnessed and promises of future more elaborate and inclusive finance.

3 Evolution of Virtual Savings and Credit Technological Platform in Kenya

In the last decade, the mobile phone-based financial services platform has shaped the development of financial markets in the East African region and is now being replicated in other African regions. In March 2007, following a pilot project, the M-Pesa technological platform was launched as a bank product in partnership with a telecommunications company. This innovative development has since undergone four stages that include: the introduction of a retail electronic payments platform; pushing the financial inclusion frontier through virtual savings; integrating credit platforms; and finally pushing the frontier of cross-border remittances and electronic payments. Ndung'u (2018) shows these developments in the following four stages of digital evolution in Kenya.

The **first stage** is where the mobile phone technological platform was used for electronic money transfers between users (Person to Person (P2P), Person to Government (P2G), and Government to Person (G2P)), and later for retail electronic payments and settlement. The emergence of a retail electronic payments system was a game changer—it has been described as efficient,

effective, transparent, and safe, among the many financial services that have developed as offshoot to the original payments platform.

In the **second stage**, the platform integrated with more commercial banks, SACCOs and insurance companies, and introduced virtual savings accounts. The virtual savings accounts developed into a technological platform for managing micro-deposits and micro-savings accounts, payment of utility bills, insurance premiums to insurance companies, and settlement of insurance claims. This made it easier for banks to provide accounts for the unbanked and pushed financial inclusion frontier. A new definition of digital financial services emerged, but more importantly, it improved the banking intermediation process. Consequently, commercial banks in Kenya built mountains of deposits and expanded their networks across the country and the East African region. In 2009, Safaricom launched its pay bill services. In the subsequent expansion, Safaricom partnered with 25 banks and over 700 businesses to facilitate fund deposits, bank transfers, and the regular payment of utility bills, insurance premiums, and loan installments. Other mobile network operators (MNOs) followed with similar products and competitive services. In addition, others moved to mobile virtual network operations (MVNOs)—a wireless communications services provider that does not own the wireless network infrastructure over which it provides services to its customers. Businesses using the online shopping platform to reach customers who eventually make their payments through the digital finance platform have expanded. Additionally, most utility companies have embraced electronic payments. Consequently, sustainable business models have been developed on the digital platform across all sectors of the economy.

The **third stage** is where the mobile phone financial services platform developed from virtual savings account to virtual credit supply platform. The M-Shwari, launched in November 2012 through a strategic partnership between Commercial Bank of Africa (CBA) and Safaricom, became the virtual platform in Kenya at this stage. After 40 months of operations, M-Shwari had over 15 million customers. It has now expanded to Tanzania as M-Pawa and to Uganda and Rwanda as MoKash and to Cote d'Ivoire as Momo-Kash. The novelty of this stage is the use of transactions and savings data to generate credit scores used to price micro credit and assess credit risks. Other products were launched and have been replicated across the East African region. In March 2015, Safaricom partnered with Kenya Commercial Bank (KCB) to rollout KCB M-Pesa—a mobile phone-based savings, credit supply and retail payments transactions platform. In addition, other products such as M-Kesho, Tangaza, and Mobicash have led to increased access to credit and improved savings in the country. Transactions on the virtual

platforms have generally become the entry point for financial services, with the informal markets also using these formal financial services. The **fourth stage** is where cross-border financial transactions developed and using the same mobile phone-based transactions model.

Following the adoption of virtual savings and credit products, savings culture has emerged as evidenced by the growth in the micro savings. For instance, in Kenya, the digital revolution in the financial sector has provided a means of managing bank accounts enabling commercial banks to reach out to more customers and grow deposits that provide them with strong intermediation capacity. But the critical issue here is that banks were provided with an easy technological platform to manage micro-accounts. Figure 1 shows the growth of micro-accounts and gross deposits in Kenya since 2005. This is the macro level outcome that is being pushed by the developments at the macro and micro levels with regard to how alternative finance is working.

Figure 1 shows that the number of micro-accounts have increased from 2.14 million in 2005 to over 50 million as at April 2018. The data also show that the deposit accounts with amounts over KSh 100,000 (approximately US\$1000) registered a slight increase over the same period increasing from 0.41 million in 2005 to 1.62 million as at the end of 2017. The growth in deposit accounts, dominated by the micro deposit accounts, has been accompanied by an increase in gross deposits from \$6868.0 million in 2005 to \$31,170.15 million in April 2018. These data points show that access to financial services and enhanced financial inclusion seem to have taken place in Kenya.

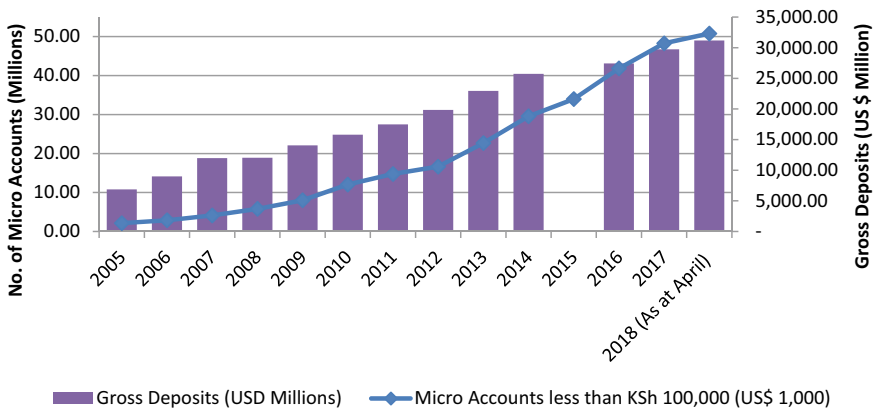


Fig. 1 Growth in micro accounts and gross deposits in Kenya (Source Central Bank of Kenya, 2018)

4 Digital Evolution Pushing Frontier of Financial Inclusion in Africa

Evidence is emerging that the success of financial inclusion, especially for women and low-income earners, in African economies is being driven by the success of retail electronic payments system. Alternative finance story seems to revolve around the binding constraints on financial services that digital evolution has helped to overcome. Among the binding constraints that sustained financial exclusion in Africa were levels of income, the irregular flow of that income and the physical distance to a bank branch or financial service point. In one stroke, the digital financial services seem to have solved these binding constraints and turned tables from financial exclusion to financial inclusion. Most participants who entered this ecosystem did not have or did not need a bank account. It was easy access for them as the mobile phone number became the account number. For those with bank accounts, they were integrated with their bank accounts and could withdraw money directly from their accounts using their mobile phones. All mobile phone account holders would then benefit from a range of financial services. These have pushed the frontier of financial inclusion in Africa, as evidenced by the Kenyan case shown in Fig. 2. But more importantly it has enabled the retail electronic payments that are the entry point for financial services.

Kenya's financial inclusion has tripled in just 13 years with 82.9% of the population now formally included. Figure 2 shows that access to formal financial solutions has improved from 26.7% in 2006 to 82.9% in 2019. Formal financial solutions include banks, SACCOs, Micro-finance institutions, insurance, mobile money, asset finance, and digital credit applications. However, a number of people still use the informal financial solutions as alternatives on top of the formal ones. Informal financial solutions include

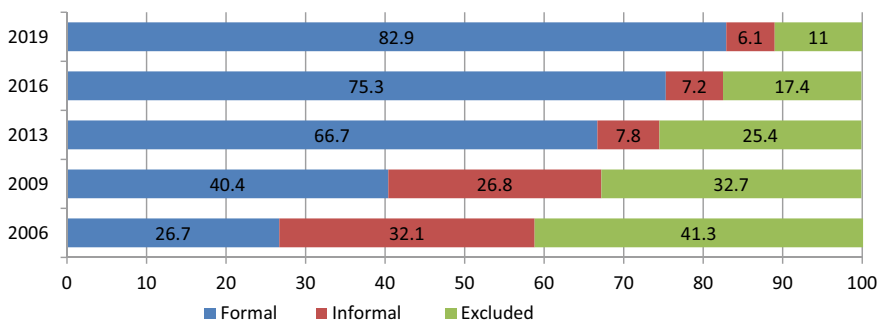


Fig. 2 Access to financial services and products in Kenya (Source FinAccess Survey 2019)

chamas (groups), moneylenders, shopkeeper's goods/cash on credit, finance from social networks, and savings at home. On the other hand, adults who are excluded from Kenya's financial services have reduced from 41.3% in 2006 to 11% in 2019. The tremendous drop has been the preference of informal financial services from 32.1% in 2006 before M-Pesa to 26.8% in 2009 after M-Pesa and drastic reduction to 6.1% of the adult population in 2019.

Since 2012 when M-Shwari demonstrated the viability moving from a transactions platform to a virtual savings and credit supply platform, the partnerships between Mobile Network Operators (MNOs) and banks have multiplied. This has led to development of a number of digital savings and credit products in Kenya. These include: M-Co-op cash by Co-operative Bank of Kenya, KCB M-Pesa by Kenya Commercial Bank, and Timiza by Barclays bank (now ABSA Bank). In July 2015, Equity Bank's subsidiary Mobile Virtual Network Operator (MVNO)³ launched its product that permits its customers to use its new mobile phone service (Equitel) with other mobile network operator services. Thereafter, in 2017, Equity Bank rolled Eazzy loan, a digital savings and credit products. Other unregulated application-based Fintech's products such as Tala and Branch have developed in Kenya. The development of similar digital financial products has been replicated across Africa. The Kenya case allows for the conclusion to be made that millions of adults now have the possibility of borrowing using their phones, which has opened the door to private, formal consumer credit for the first time (Gubbins and Totolo 2018). Evidently, the FSD Kenya (2019) report showed that over 70% of the Kenyan adult population is saving formally or informally whereas 50% of the population currently has a loan product. MicroSave Consulting (2019) pointed out that the number of digital loans issued from 2016 to 2018 has approximately doubled, with about 86% of the loans taken being digital—digital loans issued increased by 1.9 times in 2018. The study also noted that despite the entry of almost 50 fintechs in the four years to 2018, bank and MNO-facilitated products dominated the market, amounting to 97% of supply. Figure 3 shows the trend in savings and credit products in the country.

Figure 3 shows that following the launch of virtual savings and credit products (led by launch of M-Shwari in 2012), the savings prevalence increased from 66% in 2013 to 80% in 2016. However, this declined by ten percentage points to 70% between 2016 and 2019. The decline in savings prevalence

³ MNO refers to mobile network operators who buys radio spectrum and provides end-to-end service (voice or data) to its subscribers. On the other hand, MVNOs buy voice and data packages in bulk from MNOs and sell them to their subscribers/users. MVNOs do not have government license to use their own radio frequency.

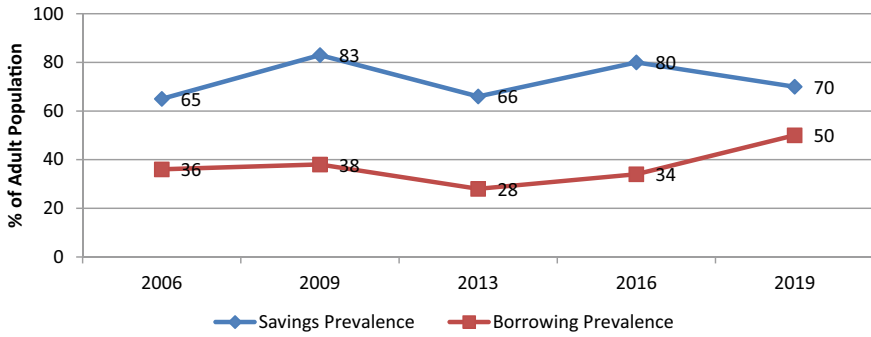


Fig. 3 Prevalence of savings and credit products in Kenya (% of Adult Population) (Source FSD Kenya 2019)

is related to recent economic downturn in Kenya. This may have pushed individuals to over borrow based on their previous virtual savings and credit history as evidenced by the borrowing prevalence that has been on a steady rise from 28% in 2013 to 50% of the adult population in 2019. Figure 4 shows the saving solutions in Kenya’s formal digital savings products relative to that of formal non-digital and informal savings products.

Figure 4 shows overlapping usage of the savings solutions by the individuals included in the sample. Figure 4 shows an increase in savings using the formal digital solutions from 11% in 2009 to 56% in 2019, an indication that digital platforms have facilitated savings growth and credit uptake. On the other hand, savings using formal non-digital solutions declined from 23% in 2016 to 18% in 2019. According to FSD Kenya (2019), between 2016 and 2019, use of mobile phone savings products in Kenya increased from 43% to 54% while mobile banking remained at 17% in 2016 and 2019. However, use of SACCO savings products declined from 13% in 2016 to 11% in 2019. Similarly, use of bank saving products (excluding mobile banking products)

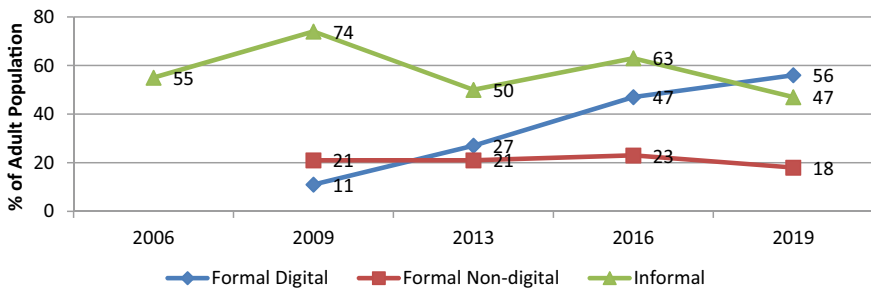


Fig. 4 Savings solutions in Kenya (% of Adult Population) (Source FSD Kenya 2019)

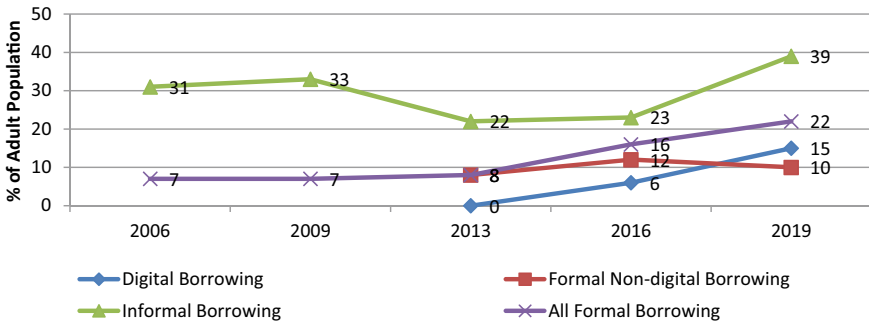


Fig. 5 Borrowing trends in Kenya (% of Adult Population) (Source FSD Kenya 2019)

declined from 10% in 2016 to 8% in 2019. Savings using informal solutions (including secret hiding place, and friends and family) also declined from a peak of 74% in 2009 to 47% in 2019. This shows that the increase in formal savings in the country can largely be attributed to the uptake of the virtual savings products.

Figure 5 shows the trend of virtual credits since 2013 in comparison to other formal and informal financial solutions (such as shopkeeper credit, shylocks, family, and friends).

Figure 5 shows that both formal and informal borrowing has risen between 2013 and 2019. More importantly, Fig. 5 shows that digital borrowing rose from 0% in 2013 to 15% 2019, pushing all formal borrowing from 8% in 2013 to 22% in 2019. FSD Kenya (2019) noted that the overall formal borrowing figure of 22% is lower than formal digital and formal traditional combined because of overlapping usage, that is, a borrower may have both a formal digital and formal traditional loan simultaneously. Specifically, FSD Kenya (2019) showed that the population using virtual credits rose from about 0.6% in 2016 to 7% in 2019 while those using mobile banking loans increased from 6% to 9% over the same period. MicroSave Consulting (2019) noted that the total value of traditional credit from banks is still higher than virtual credit, but the number of virtual loans has increased substantially across all provider categories. The study also showed that by 2018, about 77% of borrowers had taken only virtual loans, 10.2% had taken only traditional loans, and 12.4% had taken both virtual and traditional loans. This reflects the wider accessibility and reach of virtual credit for borrowers in Kenya. The FSD Kenya (2019) report also pointed out that the above trends suggested that the virtual credits are not substituting for traditional sources of credit, but instead are reaching a new audience of borrowers—perhaps people who were not satisfied with existing options or previously unable to obtain credit at the size and terms they needed.

5 The Virtual Savings and Credit Supply Products Developing in Africa

The development of virtual savings and credit products that began with M-Shwari in Kenya has been replicated across Africa. The similar products developed include KCB M-Pesa, M-Co-op Cash, Eazzy loan and Timiza in Kenya; M-Pawa in Tanzania; Mokash in Uganda and Rwanda; and MoMoKash in Côte d'Ivoire. This section provides evidence on the changing structure of financial services and the future of banking in Africa.

5.1 M-Shwari in Kenya

M-Shwari, a mobile-centric banking solution offered by Commercial Bank of Africa (CBA) in partnership with Safaricom, enables a customer to open and operate a virtual savings accounts if they already have an M-Pesa account. It means a customer can activate a virtual savings account and transfer savings from M-Pesa to M-Shwari without any charges. A customer can save as low as one Kenya Shilling and earn interest of up to 6.65% per annum on the savings balance. Moreover, it enables a customer to access micro-credit of a minimum of KSh. 100 (approximately US\$1) any time, which is received instantly through the M-Pesa account charged at a facility fee of 7.5%. If a borrower defaults, the repayment period is extended by additional 30 days and an additional 7.5% facilitation fee is applied. Thereafter, the borrower loses access to M-Shwari for 30 days for late loan repayment. M-Shwari also offers lock savings account—a flexible fixed deposit account that allows savers to customize their savings depending on the purpose of their savings. The lock savings account earns interest of up to 70% of the Central Bank rate. Figure 6 shows the adoption of M-Shwari up to June 2019.

As at June 2019, M-Shwari had 28.8 million customers in 6.5 years of operation. Out of these, 56% or 16.1 million were males and 12.7 million or 44% were female.

5.2 M-Pawa in Tanzania

M-Pawa, coined from an English word that means to “enable or to facilitate,” is a mobile-based banking service offered to Tanzania M-Pesa customers by Commercial Bank of Africa (CBA) in partnership with Vodacom. M-Pawa was launched in May 2014. It allows customers to save money through their

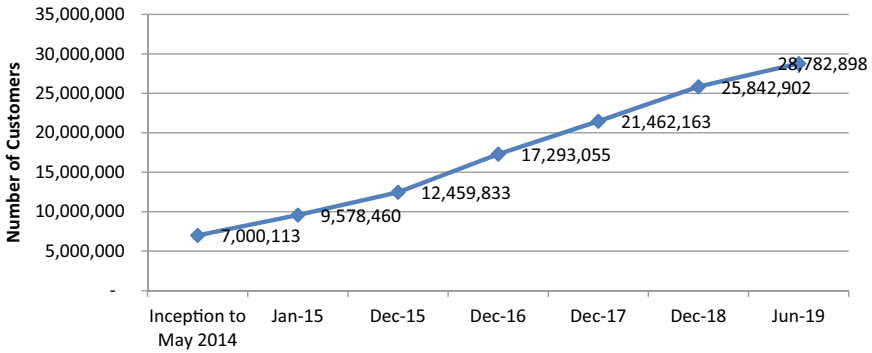


Fig. 6 Growth in M-Shwari customers in Kenya (Source Commercial Bank of Africa, 2019)

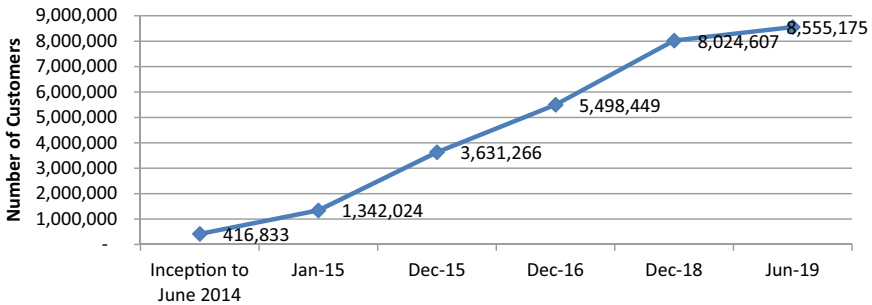


Fig. 7 Growth in M-Pawa customers in Tanzania (Source Commercial Bank of Africa, 2019)

mobile phones, earn interest from savings, and access micro-loans. The M-Pawa deposit account is a micro-savings product that earns up to 5% interest with no charges levied on the account, and no charges for moving money between M-Pesa and M-Pawa accounts. Figure 7 shows the growth in M-Pawa customers since inception.

Figure 7 shows an accelerated growth in M-Pawa customers in the period January 2015 to December 2018, and to about 8.5 million in June 2019.

5.3 MoKash in Uganda and Rwanda

MoKash, coined from an English phrase “More Cash,” is a mobile-based banking service offered to Uganda and Rwanda MTN mobile money customers by Commercial Bank of Africa (CBA) in partnership with MTN. MoKash was launched in Uganda in August 2016 and in Rwanda in February

2017. The MoKash micro-savings product earns an interest of up to 5% and has no charges on the account as well as when moving money between MTN mobile money to MoKash accounts. The MoKash loan account, a micro-credit product, levies a one-time fee of 9% with a repayment period of 30 days. If the loan is not paid within the 30-day period, customer may extend the outstanding loan for another 30 days at a fee of 9%. The customer's average MoKash savings balance and loan performance (loan uptake and repayment) affects their loan limit. If a customer does not pay the loan within 30 days, the amount is recovered from their MoKash savings account and/or their mobile money account. Any balance on the principal amount of the loan after attempted recovery attracts a roll-over fee of 9% and is added to the outstanding principal amount of the loan balance. Figures 8 and 9 show the growth in MoKash customers in Uganda and Rwanda, respectively.

The figure has a few data points, but they show a steady growth in the number of MoKash customers in Uganda from 185,100 in August 2016,

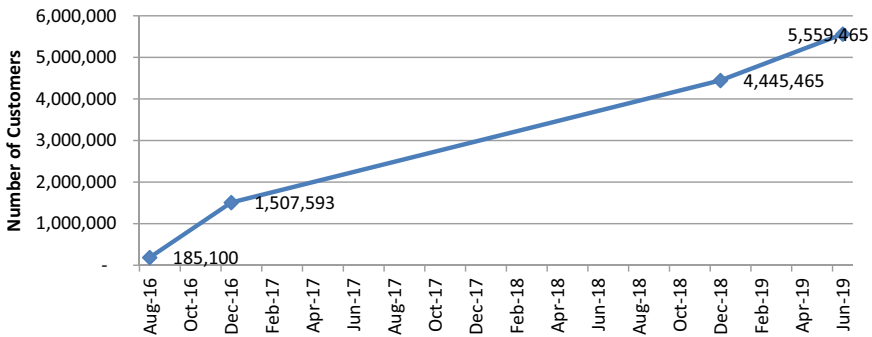


Fig. 8 Growth in MoKash customers in Uganda (Source Commercial Bank of Africa, 2019)

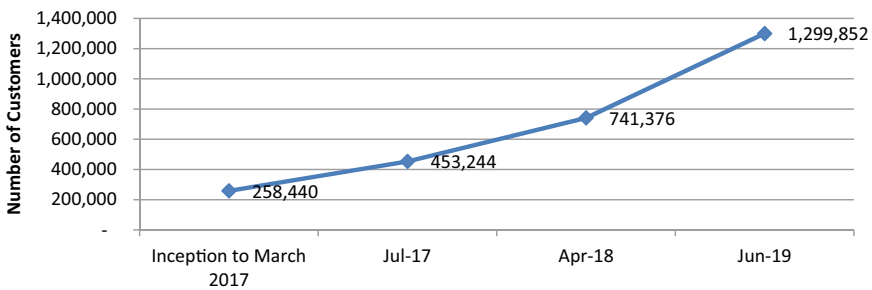


Fig. 9 Growth in MoKash customers in Rwanda (Source Commercial Bank of Africa, 2019)

when the digital product was launched, to more than 5.5 million customers in June 2019. As of June 2019, about 6189 customers used MoKash products per day in Uganda, with over 2 million active customers on aggregate. The non-performing loan (NPL) ratio was at 5.73% for the provisions against disbursement in June 2019.

As shown in Fig. 9, there has been an exponential growth in MoKash customers in Rwanda since inception in February 2017. As of June 2019, the number of MoKash customers in Rwanda stood at about 1.3 million and about 1399 customers used the MoKash products in Rwanda per day. The non-performing loan (NPL) ratio for the scored customers was at 4.04% in June 2019.

5.4 MoMoKash in Ivory Coast

MoMoKash is the first micro-loan and micro-savings product offered on the mobile money platform in Cote d'Ivoire. The Commercial Bank of Africa (CBA), through Bridge Bank West Africa, together with MTN Ivory Coast developed the virtual savings and credit supply product, which was launched in January 2018. Through the MoMoKash deposit account, one can save as a little as 50 CFA francs and up to 10 million CFA francs. The deposit account does not require a maintenance fee and attracts an interest on savings ranging from 3.5% up to 7% per year, paid out every 3 months. Through the MoMoKash loan account, one can borrow a micro-loan ranging from 2000 CFA francs to 100,000 CFA francs, paid over a period of one month. Figure 10 shows the growth in adoption of the MoMoKash product in Ivory Coast since inception.

Figure 10 shows an accelerated growth in number of MoMoKash customers between June 2018 and June 2019. This trend is expected to

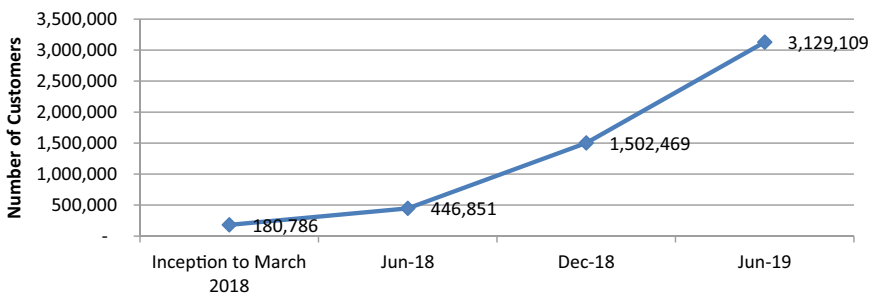


Fig. 10 Growth in MoMoKash customers in Cote d'Ivoire (Source Commercial Bank of Africa, 2019)

continue as mobile money penetration in Cote d'Ivoire increases. As of June 2019, about 9037 customers used MoMoKash products and services per day, while over 2.1 million customers were active on aggregate. The non-performing loan (NPL) ratio for the scored customers was at 23.25% in June 2019.

Table 1 provides a summary of some characteristics of the virtual savings and credit products in the five countries, described above.

Table 1 shows when the products were launched, the number of customers as at June 2019, the interest paid on savings, loan qualifications and limits, repayment duration, facilitation fees among other characteristics. Table 1 shows the similarity of the products and the characteristics across the five countries. There is a general increase in digital savings and deposits in all the five countries. Table 1 also shows an increase in total loan amount disbursed in all the countries. This perhaps reflects the important role that the virtual credit supply products are playing as alternative sources of finance in Africa. The average usage of the virtual savings and credit platforms in the five countries show that these are not dormant accounts as they are actively used by the customers. Figure 11 shows M-Shwari's average savings and borrowings per month as well as frequency of borrowing in a year for the period 2013 to 2018.

Figure 11 shows that the average borrowings per month have been increasing while average savings per month seem to have stagnated at around KSh 546 (approximately US\$5.5) since 2014. The average borrowings per month, on the other hand, increased from KSh 1059 (approximately US\$10.5) in 2013 to KSh 3652 (approximately US\$36.5) in 2018. The average savings and borrowings show that the virtual savings and credit products are mainly by low-income earners. The frequency of borrowing within a given year has been on an increase during the period. This rose from five borrowings per year in 2013 to 16 borrowings per year in 2018. This is an indication of possible loan stacking (where a loan or cash advance is approved on top of an existing loan or advance that is already in place with similar characteristics and payback terms) though most virtual credit providers require their customers to clear outstanding loans first before additional credit is advanced. This affirms that the virtual savings and credit accounts are not dormant accounts. This is further illustrated by the growth in retained deposits and loans disbursed as shown in Fig. 12.

Figure 12 shows that M-Shwari's retained deposits increased from KSh 2232.8 million (approximately US\$22.3 million) in 2013 to KSh 15,090.2 million (approximately US\$15.1 million) in 2018. On the other hand, the total loans disbursed rose from KSh 7149.8 million (approximately US\$71.5

Table 1 Characteristics of the virtual savings and credit products

	M-Shwari (Kenya)	M-Pawa (Tanzania)	MoKash (Uganda)	MoKash (Rwanda)	MoMoKash (Ivory Coast)
Launched	November 2012	May 2014	August 2016	February 2017.	January 2018
Number of Customer (as at June 2019)	28.8 million	8.6 million	5.5 million	1.3 million	2.1 million
Interest on Savings (per annum, up to)	6.65	5	5	5	7
Loan Qualification	An active M-pesa account for at least six months	An active M-Pesa account for at least 6 months, save on M-Pawa and actively use M-Pawa services	MTN mobile money subscriber at least 6 months, save on MoKash and actively use other MTN services	MTN mobile money subscriber at least 6 months, save on MoKash and actively use other MTN services	MTN MoMo account, be well identified and be an active MoMoKash user
Loan Limit	Based on the M-Pesa transactions history, savings, and past loan repayments	-	UGX 1,000,000	Rwf 300,000	100,000 CFA francs
Repayment Duration	30 days	30 days	30 days	30 days	30 days
Facilitation Fee	7.5%	9%	9%	9%	-
Total Deposits as at June 2018 (US\$ Million)	140.20	8.04	1.48	0.51	-
Total Deposits as at June 2019 (US\$ Million)	166.60	7.79	1.98	1.06	7.84

(continued)

Table 1 (continued)

	M-Shwari (Kenya)	M-Pawa (Tanzania)	MoKash (Uganda)	MoKash (Rwanda)	MoMoKash (Ivory Coast)
Average Savings as at June 2018 (\$)	5.97	1.09	-	0.65	-
Average Savings as at June 2019 (\$)	5.82	0.91	0.36	0.81	2.51
Total Loan Amount Disbursed as of June 2018 (US\$ Million)	2771.63	56.43	19.81	4.11	-
Total Loan Amount Disbursed as at June 2019 (US\$ Million)	3749.71	69.52	50.27	13.46	6.87
Average Usage per day (as at June 2019)	16,333 customers	2953 customers	6189 customers	1399 customers	9037 customers

Source Commercial Bank of Africa, 2019

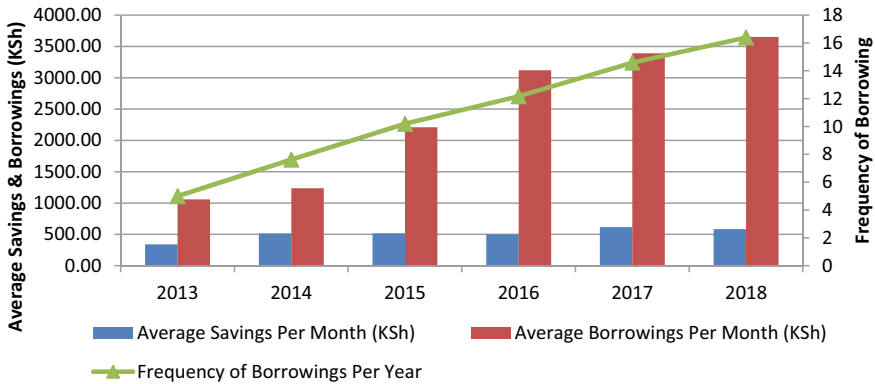


Fig. 11 M-Shwari average savings and borrowings (Source Commercial Bank of Africa, 2019)

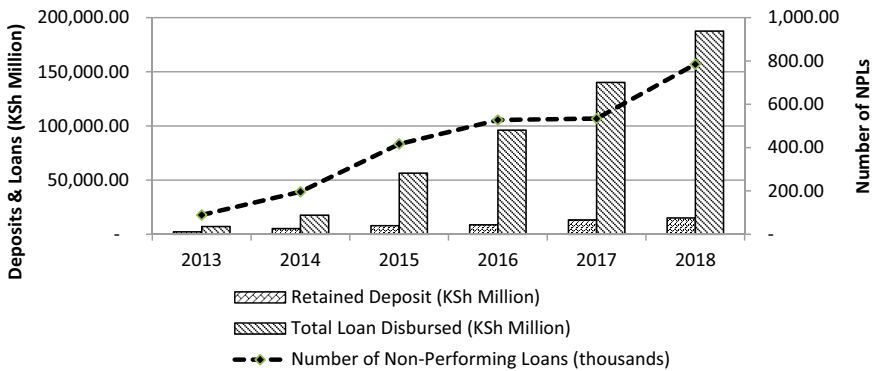


Fig. 12 M-Shwari’s retained deposits, loans disbursed and non-performing loans (NPLs) (Source Commercial Bank of Africa, 2019)

million) in 2013 to KSh 187,599.8 million (approximately US\$187.6 million) in 2018. The number of non-performing loans (NPLs) has also increased over the years with the increased borrowings. MicroSave Consulting (2019) pointed out that about half (49%) of the digital credit borrowers with non-performing loans had outstanding balances of less than US\$10. However, the non-performing loans were declining, with a 15% decrease between 2016 and 2018. MicroSave Consulting (2019) argued that this could be due to validation of the theory that overtime virtual credit performance will improve as customers “test and learn” and overtime digital loans are also increasingly valued, which in turn motivates good repayment behavior. This could also be related to the fact that regulations have been tightened and digital lenders have formed an association to contain non-performing loans.

6 Who Participates in This Market Segment?

Available data on the virtual banking platforms show some of the characteristics of the consumers of the virtual savings and credit products as at June 2019. For M-Shwari, women are using the savings product more than men. The men over 40 years are saving at a higher rate than those below this age—consistent with the life cycle hypothesis. Generally, those aged 25–34 years old borrow more across all genders, with the males contributing to about 71% to this value. The age composition of M-Shwari customers as at June 2019 was as shown in Fig. 13.

Figure 13 shows that 14% of the M-Shwari customers are aged between 18–24 years, 37% between the age of 25–34 years and 32% between 35–55 years and only 17% above 55 years of age. So this is a platform for the young: 51% of the adults below 35 years of age. These are not dormant accounts: 16,333 customers apply for loans every day and the average loan size is US\$38.27 for 29 days.

The study by MicroSave Consulting (2019) also showed that the women in the sample utilized the savings products more. Males over 40 years old seem to save at a higher rate. Borrowers are predominantly male and under 35 years old. The age cohort of 25–34 years hold the large share of outstanding loans and males hold 71% of the loan portfolio. The relative risk aversion of women borrowers is evidenced by their lower uptake of virtual credit. MicroSave Consulting (2019) noted that most women who borrow do so for household and business needs and on average take smaller loans than men. The default rate among women above 55 years was found to be higher than that of men in the same age set. MicroSave Consulting (2019) argue that this could be an indication that older women struggle more with digital literacy and with using the virtual products; or possibly the other family members use their

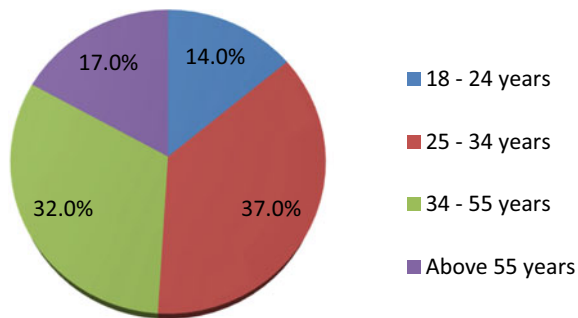


Fig. 13 Age composition of M-Shwari customers as at June 2019 (Source Commercial Bank of Africa, 2019)

IDs to access virtual loans. They should be more risk averse and less prone to indebtedness.

For M-Pawa in Tanzania, both urban centers and rural dwellings have witnessed the same level of uptake and activity. The M-Pawa savings and borrowings tend to peak at the start and end of each week while the value of retained customer deposits remain higher than the credit sought. MoMoKash product in Cote d'Ivoire mainly targeted the unbanked (who need easy access to credit or a secure savings account), retailers, and artisans. MoMoKash product also targeted students over 21 years old (who need a financial boost to meet daily needs) and parents. The available data show that over a third of the adults in these countries use virtual credit and that individuals in rural areas have not been left behind in access to digital credits but also represent the largest addressable market (individuals with mobile phones who have not used digital loans).

Evidently, the developments in alternative finance have benefited the unbanked (mostly women) and the young population that is technologically savvy. Both urban and rural dwellers, retailers, and artisans (who earn small daily incomes and seek a safe and profitable way to save and borrow money in order to expand their business), and poor households (to deal with urgent situations of their families) have benefitted. However, the digital divide is still a concern in the African countries. For instance, internet connectivity is still less reliable in rural areas and access to internet is still costly to the low-income households. MicroSave Consulting (2019) pointed out that factors, such as the quality of smartphones, data coverage, high mobile data price sensitivity among low-income users, and to some extent lower digital literacy for women and rural customers further risks excluding these consumers who are already underserved.

7 The Emergence of Fintechs that Are not Regulated as Digital Bank Products

In view of the increasing demand for short-term micro-credit, application-based (App-based) fintech lenders have emerged. However, they are not regulated and are not mandated to report to the credit reference bureaus like the Telco-facilitated bank virtual savings and credit products. Several startups (such as *Tala* and *Branch*) offering credit through Apps have emerged. *Tala* was the first instant mobile loans app launched in Kenya as *Mkopo Rahisi* (Easy Loan) in 2014. To access *Tala* loans, a customer needs to have an android smartphone as well as an active Facebook account. The App mainly

depends on reading the borrower's M-Pesa transaction history to determine whether they qualify for the loans. *Branch* also works exclusively through a mobile application thus requires customers to have an android phone to make it easier for the App to gain access to M-Pesa transaction messages. The M-Pesa transaction messages are used to determine a borrower's eligibility and loan limit. More unregulated virtual credit supply platforms have mushroomed in Kenya. These include *Saida*, *Haraka*, *Okash*, *Pesa Pata*, *Pesa na Pesa*, *Zidisha*, and *Kiva*. Most of them borrowed the *Branch* and *Tala* business model. They are mobile Apps that rely on M-Pesa or Airtel money transactions to determine eligibility and quantum of loans. Some, like *Zidisha* and *Kiva*, are peer-to-peer (P2P) credit platforms, lending money to individuals or businesses through online services that match lenders directly with borrowers. The mushrooming of such unregulated virtual credit only platforms could be an indicator of regulatory failure in the African countries. Table 2 presents some statistics on a sample of regulated and the unregulated players in Kenya.

Table 2 shows that the unregulated App-based fintech lenders serve a significant market segment. MicroSave Consulting (2019) showed that the leading App-based lenders account for about 3% of the market. The main challenge in the market is that there are no fintech regulations in place. Though the Capital Market Authority launched a sandbox for fintechs, the digital lenders are currently not included. There are no regulations on lending limits and operational standards. In March 2019, about 12 fintechs formed Digital Lenders' Association of Kenya (DLAK), with a code of conduct to promote industry best practices and to drive a coordinated approach in addressing their pressing issues.

The steady growth in adoption of the virtual savings and credit supply products and the number of customers that utilize the products per day mirror the growth in micro-accounts and the active accounts, which show increased micro-savings and micro-borrowings in the five African countries. Indeed, the virtual savings and credit products have become reliable and inclusive alternative source of financial services and short-term credit to many in the African economies, where the development of the credit market, collateral technology has been quite constrained and a binding constraint to finance.

Table 2 Some regulated digital lenders and unregulated app-based fintech lenders in Kenya as at October 2018

Players	M-Shwari	KCB M-PESA	MCo-op cash	Eazzy loan	Tala	Branch
Regulation Summary of model	Regulated Telco-facilitated bank			Bank + Mobile virtual network operator (MVNO)	Unregulated Fintech (Android app focused)	
Total subscribers (million)	20.10	9.80	3.30	1.60	0.77	0.75
No. of loans disbursed (million)	83.30	15.40	2.80	4.20	1.80	1.50
Value of loans disbursed (million USD)	2080	482	87	570	35	20
Approximate daily loans	70,000	21,000	1000	8500	310	190
Loan portfolio (million USD)	80.0	24.0	8.6	38.0	7.8	4.0
Loan interest rate (monthly)						
Loan interest rate (monthly) (%)	7.5	3.66	3.66	3.66	15	1–14
Loan size range (USD)	1–500		10–1000		20–500	2.5–500

Source MicroSave Consulting (2019)

8 The Digital Platform for Alternative Finance—Some Conclusions

In this section, we have made a case of alternative finance that is developing, pushed to the frontier by the digital technology. We have used a

case for virtual banking services with interesting characteristics. We have shown that effective, efficient, transparent, and safe retail electronic payments platforms have emerged creating an easier entry point for financial services and formalizing the informal market. The alternative finance platforms are largely operated from commercial banks, but in future, they can be stand-alone institutions. Savings culture has emerged as evidenced by the growth in virtual micro-savings. The virtual banking services platforms have provided commercial banks a technological platform to manage micro bank accounts for small savers enabling commercial banks to reach more customers and grow mountains of deposits. Subsequently, the huge deposits have provided banks with capacity to grow and intermediate. This has led to the emergence of strong banks leveraging the virtual platform to manage these micro-accounts, build deposits, and extend financial services to previously unbanked and undeserved population. This has improved the space for formulating an effective and forward-looking monetary policy and allowed improvements in the AML/CFT regime in some of these countries like Kenya and Tanzania.

Financial inclusion and financial development are now evident in those countries that have embraced digital financial services. Recent data show that most Africa economies are improving their financial inclusion strands as well as supporting strong banks with huge deposits that allow them future capacity to grow, intermediate, and innovate. Evidently, the digital finance platform is driving growth and development of the financial sector in Africa. FSD Kenya (2019) showed that there has been an 8% increase in the share of the population using mobile money, mobile banking, and digital apps since 2016 in contrast to a 2% decline in users of traditional services (traditional bank accounts, SACCOs, and Micro Finance Institutions (MFIs)). According to the FSD Kenya (2019) report, one in every four Kenyans now uses mobile banking services that offer interest-bearing savings and formal credit. The utilization of virtual credit supply platform has then been very exciting. For instance, before M-Shwari, there were only 13,000 open loan accounts in the hosting bank. However, with M-Shwari, the host bank opened over twice as many loan accounts between 2013 and 2015—27,400 on average per day, effectively transforming the host bank from a niche bank serving corporate clients to a bank serving the mass market (Gubbins and Totolo 2018).

The innovative virtual platforms reach poorer segments of the population that are sensitive to financial products, their price, and delivery channels, especially for African economies with segmented markets. We have seen that these virtual products have navigated across market segments. Therefore, through its effect on the development of virtual savings and credit supply platforms, it will improve both financial inclusion with a wider outreach and

a more vibrant financial market. The credit supply platforms have provided avenues for virtual savings, virtual short-term credit supply and it is now evident that they have become sustainable poverty reduction pathways as well as entry point for women and the poor into financial inclusion and formal savings platform. Women can save in platforms and products that cannot be encroached and they are increasingly able to borrow and invest. This has lifted more female headed households from poverty. A recent study by Suri and Jack (2016) showed that 2% of Kenyan households, mostly female-headed, have been lifted from poverty. The digital finance platform facilitates the confidentiality and convenience that women require in financial services and payments. These developments have been replicated across African economies where virtual banking services have been adopted as the examples in five countries in Africa portray.

Alternative finance is about the future of financial services. The cases provided here by M-Shwari in Kenya and four other countries are good examples. Branchless banking has been embraced and the service turnaround time has been greatly reduced. For instance, it takes about five seconds to save or withdraw from a virtual savings account to a transactions account. Access to credit takes the same turnaround time and the collateral technology used is transformative too. The average loan size is about \$38, for example in Kenya, payable within 30 days. In Rwanda, it is \$5–\$10 payable in 30 days. The number of active accounts has increased showing an improvement in saving culture due to the emergence of these virtual products and services. In Kenya, for example, one product (M-Shwari) has 28.8 million accounts for an adult population of about 35 million, about 82% coverage.

A lot can be said about the future of financial services in Africa with this model of alternative finance, being propelled by digital evolution. However, the major concern to policymakers in Africa is what factors are essential to sustain the digital evolution in Africa. First, development of connectivity infrastructure for mobile phones and internet such as the fiber optic network is important since financial services can be slowed by poor connectivity. Also, scalability of the digitization process is important for sustainability. Second, interoperability of retail electronic payments platforms, market conduct, and competition are challenges that require feasible solutions. The retail electronic payments were developed individually, but now they should allow an interoperable platform that will use connectivity to enlarge the market, lower unit costs, and provide room for more innovative solutions and scalability, while at the same time allowing seamless and transformative regulatory technology to emerge. There is a need to enhance MNOs and MVNOs interoperability, agent network interoperability, and agency banking interoperability across

Africa. Third, a transformative regulatory technology is required to nudge the market to the appropriate directions and not stifle innovations—regulators should not stifle innovation in the market place. Regulatory technology must move in partnership and encourage sustainable and tractable innovations. Fourth, is the improvement in the identity system (electronic-ID system) to cope with market innovations and secure the market. Kenya's success was initially supported by the presence of an ID system that created a unique identifier for KYC purposes and security. However, the identification system needs to be upgraded to a robust, inclusive, and privacy-supporting e-ID system able to provide identification services for an increasingly digitized society. Finally, state and institutional capacity should be improved to cope with market innovations and dynamism—develop institutions that regulate the market and those that protect financial market. There is a need to further strengthen state and institutional capacity to cope with regulatory challenges, provide market safeguards, and evolve a national identification system that is suitable for the next phase of an increasingly digital economy. These factors are crucial to the continued success and development of the future of financial markets in African economies.

References

- Bennett, D., G. Mungai, and P. Richardson. 2016. *Financing SME growth in Kenya*. Nairobi: Financial Sector Deepening Trust Kenya (FSDK).
- FSD Kenya. 2019. *Inclusive finance? Headline findings from FinAccess 2019*. Nairobi: Financial Sector Deepening Trust Kenya.
- Gubbins, P., and E. Totolo. 2018. *Digital credit in Kenya: Evidence from demand-side surveys*. Nairobi: Financial Sector Deepening Trust Kenya.
- MicroSave Consulting. 2019. Making digital credit truly responsible: Insights from analysis of digital credit in Kenya. Available at: <https://www.microsave.net/wp-content/uploads/2019/09/Digital-Credit-Kenya-Final-report.pdf>.
- Ndung'u, N. 2018. Next steps for the digital revolution: Inclusive growth and job creation in Kenya, The Brookings Institution Working Paper No. 20. Africa Growth Initiative, Brookings Institution, USA. Available at: https://www.brookings.edu/wp-content/uploads/2018/10/Digital-Revolution-in-Africa_Brookings_AGI_20181022.pdf.
- Ndung'u, N. 2019. Digital technology and state capacity in Kenya. CGD Policy Paper 154, Washington, DC: Center for Global Development. Available at: <https://www.cgdev.org/publication/digital-technology-and-state-capacity-kenya>.
- Suri, T., and W. Jack. 2016. The long-run poverty and gender impacts of mobile money. *Science* 354 (6317): 1288–1292. <https://doi.org/10.1126/science.aah5309>.

World Bank. 2016. *Identification for development (ID4D): Identification system analysis*. Washington DC: The World Bank.

<https://www.safaricom.co.ke/personal/m-pesa/do-more-with-m-pesa/loans-and-savings>.

<https://cba.co.tz/m-pawa/>.

<https://cbagroup.com/rwanda/mokash/>.

<https://cbagroup.com/uganda/mokash/>.

<http://www.aboukam.net/2018/08/05/momo-kash-le-1er-service-de-micro-emp-runt-et-epargne-via-mobile/>.

<http://civdigitale.com/momo-kash-catalyseur-de-linclusion-financiere-en-cote-divoire/>.



What Does Digital Money Mean for Emerging Market and Developing Economies?

Erik Feyen, Jon Frost, Harish Natarajan, and Tara Rice

1 Introduction

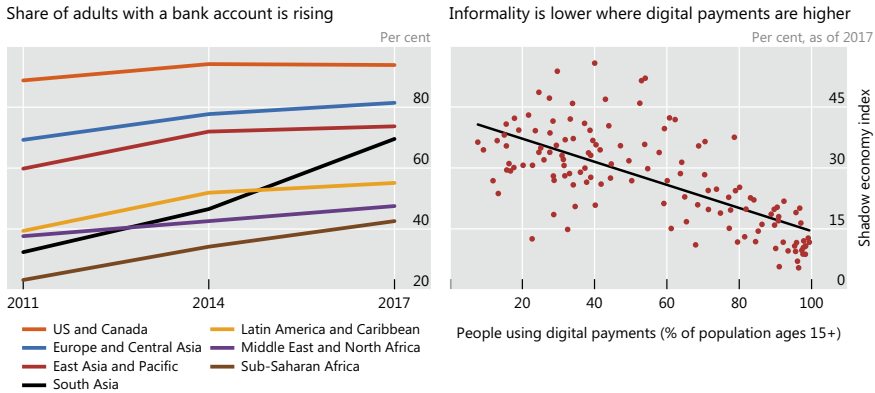
From the ancient Indian *rupya*, to cacao beans in the Aztec empire, to the first paper money in China, money and payments have been evolving for centuries. The countries that are today called emerging market and developing economies (EMDEs), which collectively make up 84% of the world's population but only 37% of GDP at current prices, are no exception. In recent decades, physical cash and claims on commercial banks (i.e. deposits) have become the main vehicles for retail payments around the world (Bech et al. 2018). Compared to physical cash, commercial bank money provides more safety, enables remote transactions and allows banks to extend other

E. Feyen · H. Natarajan
World Bank Group, Washington, DC, USA
e-mail: efeijen@worldbank.org

H. Natarajan
e-mail: hatarajan@worldbank.org

J. Frost (✉) · T. Rice
Bank for International Settlements, Basel, Switzerland
e-mail: jon.frost@bis.org

T. Rice
e-mail: tara.rice@bis.org



Source: World Bank Findex data; Medina and Schneider (2019).

Graph 1 Access to bank accounts and bank services is heterogeneous, but rising (Source World Bank Findex data; Medina and Schneider 2019)

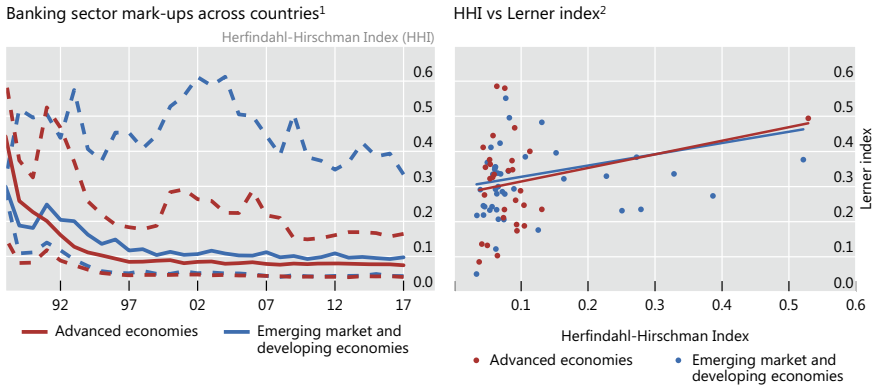
useful financial services; this may ultimately benefit economic efficiency and enhance economic policy oversight (Listfield and Montes-Negret 1994).

Yet for retail users, especially in EMDEs, commercial bank money poses at least three key challenges. First, it requires a bank account—access to which is rising (Graph 1, left panel) but is still far from universal. The poor often lack the proper documentation to comply with banks' customer due diligence (CDD) requirements. In some cases, they live too far from a bank branch, or find the maintenance costs or minimum balances too onerous. E-money, which can be seen as a variant of commercial bank money, seeks to address these challenges.¹ Together with simplified CDD and networks of agents, e-money has improved access to transaction services. Still, in countries where bank accounts and e-money have not reached universal levels, the poor rely heavily on cash. This reliance on cash helps perpetuate informality, also known as “the shadow economy”—economic activities hidden from authorities for monetary, regulatory and institutional reasons (Medina and Schneider 2019).² Indeed, informality is higher in countries with lower use of digital payments like bank accounts and e-money (Graph 1, right panel).

Second, despite improvements in recent years, financial institutions in many EMDEs face limited competition (Graph 2, left panel). This concen-

¹ E-money refers here to monetary value that is stored electronically on receipt of funds, and which is used for making payment transactions. In almost all countries, e-money balances are held in commercial banks. A notable exception is China where funds are held with the central bank.

² For a seminal work on informality, see Hart (1973). Hart described the economic activities of low-income urban workers in Accra, Ghana, including complex and varied income-generating activities operating outside the formal legal system.



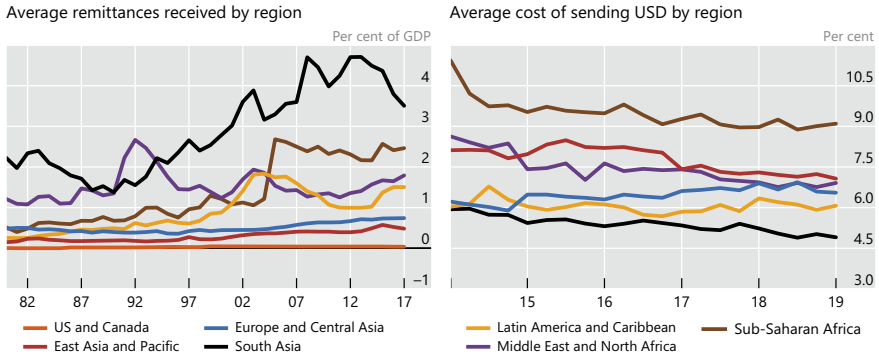
¹ Solid lines denote the median and the dash lines denote the 5th and the 95th percentiles. ² Data for 2014.
Source: World Bank.

Graph 2 Banking sector concentration, while declining, is associated with higher mark-ups (Source World Bank: Remittance prices worldwide reports)

trated market power often results in higher mark-ups (Graph 2, right panel), i.e. more expensive financial services. Concentration can also result in limited incentives for innovation over time. Together with households' recollection and past experiences of costly banking and financial crises, banking sector concentration can contribute to a lack of trust in the formal financial system.

Third, many households in EMDEs depend on low-value cross-border remittances from family members working abroad. Remittances to EMDEs reached \$551 billion in 2019. Such flows exceed official development assistance by a factor of three, and—prior to the Covid-19 pandemic—were on track to overtake foreign direct investment inflows (Ratha et al. 2019; Graph 3, left panel). Specialised money transfer operators (MTOs) have emerged to provide near instantaneous transfers, and to reduce the costs for sending money over time. Yet it still costs about \$14 on average to send \$200 back home (World Bank 2019; Graph 3, right panel). This is largely because of the need to convert remittances from and to cash on both sides of the transaction (also known as “cash-in, cash-out”). This arrangement requires manual processing (including verifying the customer's identity) and a physical office (e.g. such as an MTO or post office). Micro, small and medium-sized enterprises (MSMEs) and individuals participating in cross-border trade in EMDEs can face even higher fees and wait times than larger retail customers.

One specific problem for cross-border payments and remittances is the decline in correspondent banking. Correspondent banking is an arrangement under which one bank (correspondent) holds deposits owned by other banks (the respondents) and provides those banks with payment and other

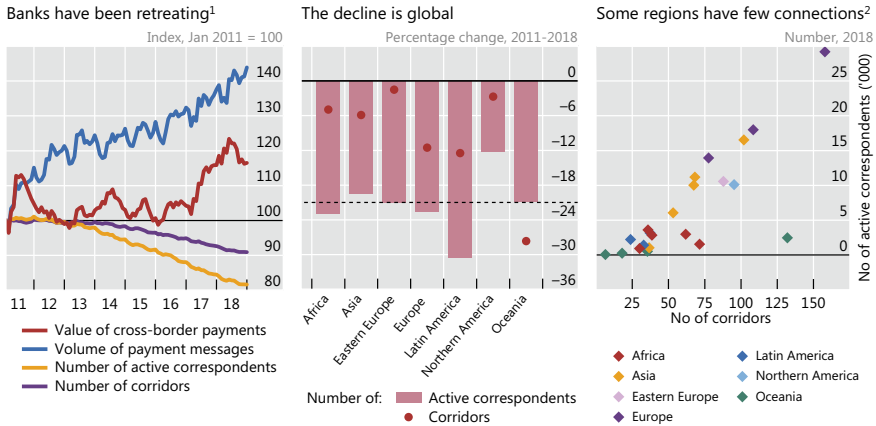


Source: World Bank: Remittance prices worldwide reports.

Graph 3 Remittance flows are increasing (Source World Bank: Remittance prices worldwide reports)

services (CPMI 2016). Most modes of cross-border payments—including banks and specialised remittance service providers—depend on the correspondent banking system, which is often slow and opaque. Moreover, in the last few years, correspondent banks have become less willing to provide such services and have been selectively exiting the business or reducing the number of respondent bank relationships (FSB 2017; IMF 2017; World Bank 2018; FSB 2019; CPMI 2019). All regions have seen a decline in the number of active correspondents, although these trends vary significantly (Graph 4, left hand panel). The rates of decline by region range from about 10 to 30%, with Northern America at the low end, and Latin America at the high end (middle panel). Additionally, the number of corridors (country-to-country connections) between countries fell by 10% over the same period. Here too, the decline was uneven across regions (Graph 4, dots in middle panel) and left some regions with fewer remaining corridors (Graph 4, right panel).

The retreat by correspondent banks raises three concerns: (1) some jurisdictions could face inadequate access to the global financial system; (2) greater concentration, or fewer correspondent banks providing services, could keep cross-border payment costs and frictions elevated and (3) where banks are not providing financial (payment) services, users may resort to less regulated or unregulated channels, shifting payments outside the banking system, including, potentially, to digital currencies (Rice et al. 2020).



¹ Three-month moving averages. ² The black dotted line shows the average percentage change of active correspondents across regions. ³ 2018 data. Averages across countries in the subregions listed. Africa = Eastern, Middle, Northern, Southern and Western; Asia = Central, Eastern, South-Eastern, Southern and Western; Eastern Europe = Northern, Southern and Western; Latin America = Caribbean, Central and South America; Northern America; Oceania = Australia and New Zealand, Melanesia, Micronesia and Polynesia.

Sources: Rice et al (2020); SWIFT BI Watch; National Bank of Belgium.

Graph 4 Correspondent banking landscape (Sources Rice et al. 2020; SWIFT BI Watch; National Bank of Belgium)

2 Enter Digital: Crypto-Assets, Stablecoins and CBDCs

Various crypto-assets claim to address deficiencies in the existing financial system. Many are vying to become a new form of digital money that can be securely sent and received over the Internet, by anybody with a phone or internet connection, and with the convenience and cost-effectiveness of an e-mail. Some initiatives target cross-border payments, particularly remittances, in EMDEs. By cutting out financial intermediaries, such proposals aim to empower users and make domestic and cross-border payments more efficient. This may be particularly relevant for country corridors hit by the decline in correspondent banking relationships, and for those countries with growing participation in the digital economy but no corresponding growth in access to e-commerce-enabled payment mechanisms.

Crypto-assets have suffered from various impediments, including high price volatility and scalability challenges, which prevent them from being adopted as a mainstream means of payment or store of value, much less a unit of account (see BIS 2018). In response, a diverse family of so-called “stablecoins” has entered the fray, including proposals like Facebook’s Libra (since renamed “Diem”). Most stablecoins attempt to maintain a stable value relative to a fiat currency (like e-money or a currency board) or a basket of

fiat currencies. To maintain a stable value, most initiatives adopt a collateral approach using bank deposits, government securities or crypto-assets although some projects attempt to maintain stability through algorithmically balancing the supply of coins in circulation with demand (Arner et al. 2020; Moin et al. 2019). This would be no small feat as the eventful history of broken currency boards and pegs has shown. Furthermore, stablecoin systems that can tap into the massive user bases of platform companies may employ network effects to drive rapid adoption on a global scale. Several big tech platform companies exist in EMDEs—in particular in Asia—that have a sufficiently large footprint to spur mass adoption.

Proposed stablecoin arrangements represent more than just a payment instrument; they are often eco-systems with entities that each play a role in the overall functioning of the system with potentially multiple digital assets that are used for payment or investment purposes running on top of them (Zetzsche et al. 2020). For most stablecoin arrangements that could reach scale, there are various key roles that are typically played by a variety of different entities:

- Governance, which includes various tasks related to software protocols, issuance and redemption policies, and the reserve investment strategy;
- Issuance and redemption of stablecoins in circulation³;
- Management of the reserve assets;
- Validation of transactions to enable transfers; and
- Custody and exchange of stablecoins with users.

However, as pointed out by the G7 and FSB, stablecoins pose a wide range of risks related to, among others, legal certainty, financial integrity, sound governance, the smooth functioning of payments, consumer protection, data privacy, tax compliance, and potentially monetary policy and financial stability (G7 Working Group on Stablecoins 2019; FSB 2020). Moreover, stablecoins face many of the same obstacles that other players have faced with transaction accounts, including mobile money. Further, they need to contend with new challenges of their own depending on the scale of adoption and their use as a means of payment or a store of value.

Recently, a number of central banks have proposed or piloted so-called central bank digital currencies (CBDCs). CBDCs would be a new form of digital central bank money that could be distinguished from reserves or settlement balances held by commercial banks at central banks (CPMI/MC

³ Some stablecoin arrangements have proposed to maintain stability by algorithmically controlling the supply of coins in circulation to match demand.

Table 1 Categorisation of cash, crypto-assets, stablecoins, bank deposits and CBDCs

	Account-based	Token-based
Central bank	Reserves/settlement balances; account-based CBDCs	Cash; token-based CBDCs
Non-central bank	Commercial bank deposits; e-money	Crypto-assets; stablecoins

2018). While the technology and design could take a number of different forms, CBDCs would be issued by the central bank, like physical cash or the reserves that banks hold at the central bank, and would be in digital form.⁴ A recent survey finds that central banks representing a fifth of the world's population say they are likely to issue a CBDCs in the next few years (Boar et al. 2020; Boar and Wehrli 2021). Several central banks are moving into more advanced stages of CBDC engagement, progressing from conceptual research towards practical experimentation. Other central banks, meanwhile, are moving at a more measured pace with further research or consultation, while yet others have concluded that the risks currently outweigh the potential benefits. Central banks give a wide range of motivations for CBDC initiatives; for EMDE central banks, this includes promoting financial inclusion and payments efficiency. Many of these initiatives target wholesale payments, i.e. large-value transactions between financial institutions, in some cases for cross-border payments. Some pilots and research and development projects are for general purpose use by retail customers (Auer et al. 2020). Economically speaking, retail CBDCs amount to households having direct access to the central bank balance sheet—"reserves for all" (Niepelt 2018).

A simple matrix helps to categorise these various digital money proposals and compare them to existing payments instruments (Table 1).⁵ The first relevant dimension is whether a payments instrument is provided by the central bank or not. In most jurisdictions, central banks play a crucial role in the payments system, holding the required reserves and settlement balances of commercial banks, and usually issuing physical cash.⁶ Private sector parties, such as commercial banks, offer bank deposits and e-money. A second dimension is whether a payments method is an "account-based" instrument

⁴ See Adrian and Mancini-Griffoli (2019) and Auer and Böhme (2020) for a discussion of different CBDC models, including models whereby a private stablecoin arrangement solely uses central bank reserves as reserve assets.

⁵ For a fuller taxonomy of money ("the money flower"), see Bech and Garratt (2017).

⁶ Notable exceptions are Hong Kong and Macau, where a limited number of commercial banks are authorised to issue bank notes for general circulation.

or is “token-based”. This distinction depends on the method of verification: the receiver of a token will verify that the token is genuine, whereas an intermediary verifies the identity of an account holder (see Kahn and Roberds 2009).⁷ Physical cash, crypto-assets and stablecoins can be considered token-based—even if the former is in physical form, and the latter are digital.

Notably, CBDCs could be either token-based or account-based depending on precise design options. A token-based CBDC would resemble a type of “digital cash”, allowing access based on knowing a password or encrypted value. An account-based CBDC would involve intermediaries like the central bank or financial institutions verifying the identity of users (Boar et al. 2020). This distinction could have some relevant implications for the use of a CBDC. For instance, a token-based CBDC could allow for greater privacy or anonymity, similar to cash. Yet this distinction may be much less important than the distinction between central bank and non-central bank instruments (Carstens 2019a; 2021).

Central bank and non-central bank payment instruments differ in a number of important ways. Crucially, the central bank is accountable to the public, rather than private shareholders. Governance frameworks have been built up over time, including in EMDEs, to safeguard central bank independence and transparency (see e.g. Crowe and Meade 2007). Forms of money offered by the central bank are provided as public goods, rather than with a profit motive. While theoretical models can assess sufficient conditions for the equivalence of public and private money (Brunnermeier and Niepelt 2019), these conditions often do not hold in practice. It is in part for these reasons that most economies, including EMDEs, feature a mix of public and private forms of money.

Indeed, central banks provide a number of key central bank public goods that underpin a stable monetary system, including providing a unit of account, guaranteeing the finality of payments, providing liquidity and conducting oversight (Carstens 2019b; BIS 2020). Those forms of money that the central bank provides—currently reserve balances and cash—usually make up only a small part of the overall money supply, but are fundamental for the functioning of the system as a whole. Meanwhile, private sector banks create money through lending by crediting a deposit account—which is steered by regulation, supervision and monetary policy. Other institutions issue e-money, which is also tightly regulated and generally kept in segregated

⁷ A second distinction raised by some authors is the degree of centralisation of the ledger. Account-based systems have a central ledger or book, while token-based systems typically run on distributed ledger technology (Bech et al. 2018).

accounts—typically in the banking system—so as to ensure safety and avoid money creation. In order to guard against excessive issuance and ensure the stability of money, substantial policy frameworks have been created, which may not yet be in place for crypto-assets and stablecoins.

Overall, digital forms of money like crypto-assets, stablecoins and CBDCs show how new technologies can be applied to address challenges in the existing monetary system, including some challenges unique to EMDEs. Yet they are to date untested at a large scale, and it is too soon to tell whether they could provide superior solutions to improving existing payment systems. Each of these innovations is evolving fast, yet understanding their risks and benefits will take time. What these innovations will mean for policy depends both on who issues them and how they are issued. In this light, the “who” may turn out to be the more important dimension. A number of practical policy challenges remain to be addressed.

3 Context in Which Digital Money May Be Adopted

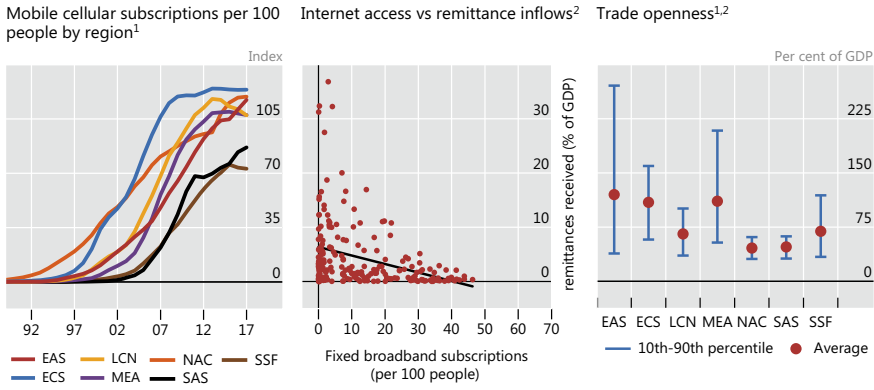
Before addressing the policy challenges in more depth, it is useful to discuss in which countries private stablecoins and CBDCs could be adopted. This is necessarily speculative, as many stablecoin arrangements and CBDCs are proposals at this point and, thus, not yet operational. Even for those projects that are live, there is scant information on adoption by country. Nonetheless, some commentators (e.g. Hileman 2015; Brosens and Cocuzzo 2019; Auer et al. 2020) have sought to sketch where crypto-assets, stablecoins or CBDCs may be attractive. We provide a similar overview, based on the discussion above, and highlight a number of relevant indicators. We discuss the potential for adoption for both stablecoins and for CBDCs, and review where the motivation for adoption could differ between them. Potential factors relate to both supply side (the digital money provider) and the demand side (the end user; household or business). These factors would also be driven by the attractiveness of the stablecoin or CBDC as a means of payment and store of value.

3.1 Supply Factors

A number of supply factors could help to support the adoption of stablecoins or CBDCs in EMDEs. Table 2 summarises these factors.

Table 2 Factors that may support the adoption of digital money

Supply factors	Description	Indicators
Infrastructure for adoption	Digital money requires a network and digital infrastructure, such as mobile phone coverage and retail agent networks, for adoption	Share of population with a mobile phone subscription, share of population with access to the Internet, availability of exchanges or MTOs for cash-in/cash-out
Traditional payment service provider profitability and costs	Incumbent financial institutional cost structures (including compliance costs) are high, making financial institutions unattractive. Digital money providers may not be subject to the same requirements (i.e. arbitrage) or could have lower compliance costs	FATF AML/CFT high-risk designation (proxies for higher KYC and risk management costs to banks); measure of off-shore/tax havens status (higher risk); incumbent financial institution profitability; level of interchange by payment card providers
Public sector desire to improve payments and financial systems	Improvements in domestic payments efficiency, payments safety and financial inclusion, reliance on cash use	Low share of population with transaction account; high reliance on cash or very low cash usage
<i>Demand factors</i>	<i>Description</i>	<i>Indicators</i>
Cost and convenience	Cost and speed of digital currency transfer or exchange may differ from traditional (cross-border) payments with a bank or MTO	Cost of receiving remittances, current speed of receiving payments
Confidence in incumbent banking system	Trust in incumbent financial institutions could be undermined by crises and concentrated markets or monopoly power	Incidence of financial crises over recent years, concentration of banking system in local market, shadow economy
Confidence in government	Trust in the public sector, including the public's expectation of sustainable monetary and fiscal policy may support CBDCs, while financial repression and weak macro-financial policies may support private stablecoins	Trust in government index, corruption perception index (Transparency International 2020), to proxy for poor rule-of-law and higher-risk countries, controls on domestic currency
Macroeconomic factors	Poor growth and large fluctuations in the value of the domestic currency may make private alternatives more attractive to users	Growth, foreign exchange volatility, inflation, trade flows



¹ EAS = East Asia and Pacific; ECS = Europe and Central Asia; LCN = Latin America and Caribbean; MEA = Middle East and North Africa; NAC = Northern America; SAS = South Asia, SSF = Sub-Saharan Africa. ² Data for 2017. Source: World Bank.

Graph 5 Stablecoins may be supplied to a greater extent in countries with higher mobile use, internet access, remittances and trade openness (Source World Bank)

Issuance and redemption of digital currencies (either private or public) requires a network and digital infrastructure, such as mobile phone coverage, for adoption. A pre-existing network could enable wide-scale adoption and make entry into markets with such networks more attractive (Graph 5, left panel). Among EMDEs, particularly countries in East Asia, the Pacific, Latin America and the Caribbean show high mobile cellular use. On a similar note, private arrangements like stablecoins may be more willing to introduce stablecoins where, for example, access to the Internet (i.e. ability to transact via the Internet) is higher (Graph 5, middle panel).⁸ Private arrangements may also be more willing to introduce digital currencies in countries with higher remittances and greater trade openness, such as the countries in East Asia and the Pacific and the Middle East and North Africa (Graph 5, right panel), as these countries would have a readier inward supply of such new payment instruments from foreign parties.

Issuance and redemption could also be driven by profitability and cost considerations of both incumbent banks and potential entrants. Such costs include entry costs (e.g. licencing fees, costs to buy or build offices and hire employees), and regulatory compliance costs (such as with anti-money laundering requirements). Private initiatives such as stablecoins would likely also require a network of physical agent offices for “cash-in/cash-out” ability

⁸ For example, the IMF (2020) notes that global stablecoin proposals could be relatively more disruptive in economies like the Philippines with high remittance inflows and high social media usage.

because most EMDEs still are extensive users of cash. Having a pre-existing network would reduce the costs and increase the scope for adoption.

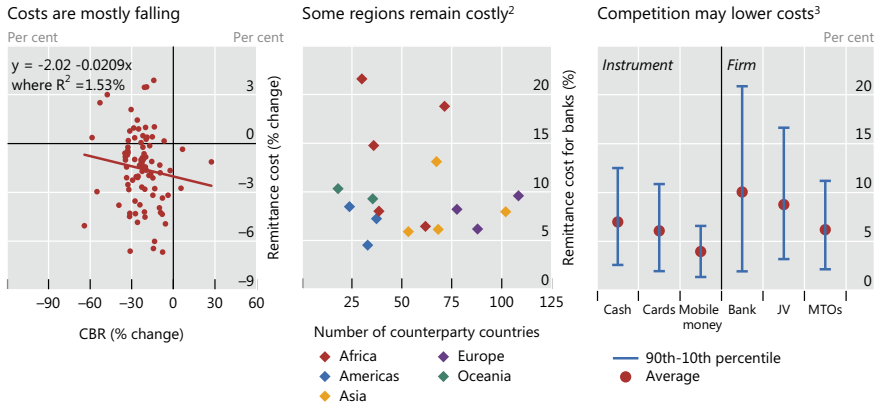
Regarding CBDCs, a number of EMDEs central banks are developing a CBDC with the aim of improving their existing payments and financial systems. Domestic payments efficiency, payments safety and financial inclusion were, on average, all considered “very important” in this respect for EMEs (Boar et al. 2020). Interestingly, a country’s reliance on cash motivated work on CBDCs but for various reasons. Those with a high reliance on cash see CBDCs as potentially reducing costs and improving know-your-customer (KYC) and anti-money laundering and combating the financing of terrorism (AML/CFT) arrangements, as set out by the Financial Action Task Force (FATF). Those with a low or declining use of cash for payments believe that a CBDC could help to maintain public access to central bank money (Boar et al. 2020).

3.2 Demand Factors

Technology is changing the way that consumers transact. They increasingly expect platforms to be mobile-first and fully digital. In EMDEs, customers look to their phones and mobile carriers to offer payment and deposit services (Petralia et al. 2019). Digital currencies have the potential to reduce the costs of transacting across borders, and increase the speed and transparency of transactions. Thus, demand for digital currencies would likely be higher in countries where those costs are relatively high, and cross-border payments are slow or opaque.

In particular, remittance costs, which have been declining over the past several years due to coordinated cross-border policy initiatives, have not declined in countries where the loss of correspondent bank access has been greatest (Graph 6, left hand panel). For some regions, particularly Africa, costs remain high (Graph 6, middle panel). Stablecoins present potentially cheaper alternatives for cross-border transactions, and would likely be desirable in countries for which receiving remittances is most expensive. Use of mobile money (discussed in detail below) and MTOs are less expensive than banks (Graph 6, right panel). Offering stablecoins through mobile phone networks and MTOs could be attractive to users and further push down the costs of cross-border payments.

CBDCs could be attractive in those countries where cash is difficult to obtain or where cash use is high, due to a lack of cash substitutes (Khiaonarong and Humphrey 2019). A number of large EMDEs, including South Africa and Mexico, show relatively high cash usage and low use of



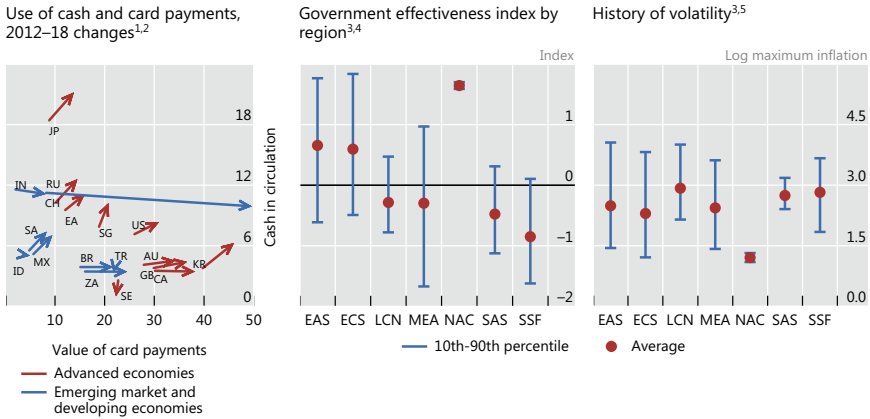
¹ The cost of sending \$200 to selected countries. ² 2018 data. Receiving subregions in each continent: Africa = Eastern, Middle, Northern, Southern and Western Africa; Americas = Caribbean, Central and South America; Asia = Eastern, South-Eastern, Southern and Western Asia; Europe = Eastern, Northern and Southern Europe. Oceania = Melanesia and Polynesia. ³ Data for 2018. Cards = credit and debit cards. JV = Joint ventures, ie partnerships between nonbank firms and financial institutions. MTOs = Money transfer operators.

Source: Rice et al (2020) using data from SWIFT BI Watch, National Bank of Belgium; World Bank, Global Findex database and Remittance Prices Worldwide, remittanceprices.worldbank.org.

Graph 6 Stablecoins could be more attractive where cost of sending remittances is high (Source Rice et al. 2020 using data from SWIFT BI Watch, National Bank of Belgium; World Bank, Global Findex database and Remittance Prices Worldwide, remittanceprices.worldbank.org)

card payments (Graph 7, left panel; Bech and Boar 2019). Countries with reduced access to banking services, e.g. due to concentration in the banking sector, may have a greater demand for CBDCs or for private stablecoins. This could also occur where there is a lack of trust in incumbent financial institutions, due, for example to a history of banking and currency crises. A relatively higher degree of financial repression (such as controls on the use of local currency or foreign exchange transactions) may make private stablecoins more attractive (Hileman 2015). Trust in the public sector, including the public’s expectation of sustainable monetary and fiscal policy may support CBDCs, while lack of effective government could make private stablecoins more attractive (Graph 7, middle panel).

Finally, macroeconomic factors may also play a role. Weak growth, large fluctuations in the value of the domestic currency or high inflation (volatility) may make private alternatives more attractive to users. This could be the case in particular for some countries in Latin America and the Caribbean, East Asia and Africa (Graph 7, right panel).



¹ AEs shown in red, EMDEs in blue. The start/end of an arrow represents 2012/2018. Data for Argentina and China are not comparable with those for other jurisdictions and are thus not shown. Data are not available for Hong Kong SAR. ² Banknotes no longer issued are not included in the calculations. For India, 2012–16 change due to demonetisation process. ³ EAS = East Asia and Pacific; ECS = Europe and Central Asia; LCN = Latin America and Caribbean; MEA = Middle East and North Africa; NAC = Northern America; SAS = South Asia, SSF = Sub-Saharan Africa. ⁴ Data for 2018. The index ranges between -2.5 (less effective) and +2.5 (more effective). It is based on 47 indicators and measures the quality of public services, civil service, policy formulation, policy implementation and credibility of the government ⁵ Data for 2017. Computed as the log of the maximum inflation rate in the past 20 years.

Source: World Bank; CPMI Red Book.

Graph 7 Adoption may depend on cash and card use, government effectiveness, volatility (Source World Bank; CPMI Red Book)

4 Particular Challenges for EMDEs

Several policy issues related to stablecoins are exacerbated in EMDEs. Authorities are confronted with six main development,⁹ macroeconomic and cross-border challenges. Table 3 provides an overview.

First, stablecoin systems could pose severe risks to the integrity of the global financial system, including for AML/CFT (FATF 2019a).¹⁰ Stablecoin systems must comply with FATF standards to mitigate their use for illicit financial activities. These standards were recently amended to cover virtual assets (VAs) and virtual asset service providers (VASPs) such as crypto-exchanges and wallets. These arrangements will now also need to conduct CDD (FATF 2019b). In their current conception, most stablecoins projects do not seek to link “accounts” to real-world identities. This raises both financial integrity and regulatory arbitrage concerns if significant volumes of transactions occur in a peer-to-peer fashion rather than using VASPs or other

⁹ Development challenges refers here to the specific policy challenges around financial sector development, including financial deepening, financial infrastructure, financial inclusion and institutional underpinnings like sound regulation and supervision.

¹⁰ Some of these risks are already apparent in the case of crypto-assets like Bitcoin. See Foley et al. (2019).

financial intermediaries. While this risk is present in all countries, authorities in EMDEs, in particular, may have more difficulty keeping pace and adjusting their surveillance, regulatory and supervisory frameworks, given resource constraints. They may also have challenges tracking and preventing financial crimes.

Second, like branchless banking and e-money networks, stablecoin systems would need to offer robust and secure “cash-in /cash-out” functions between stablecoins and fiat currency through physical agent networks since most of the local economies in EMDEs are still far from widely accepting digital payments—for mobile money such transactions accounts for about 70% of transactions (GSMA 2019). This is challenging if distribution networks are not equipped to handle crypto-asset or stablecoin transactions, lack geographical coverage or are prone to cyber-attacks. So far, it is unclear whether stablecoin systems would work on simpler “feature phones” and in locations with poor connectivity, or whether they could better address the challenges posed by a lack of ID for onboarding the unbanked, particularly in remote locations.

Third, fundamentally, stablecoins in foreign currencies or in a basket of foreign currencies will fluctuate against local currencies in EMDEs. This inhibits their adoption for daily payments since prices will remain denominated in local currencies in all but the most extreme cases. If used for debt contracts, this is a new form of foreign exchange (FX) lending. FX lending has been at the heart of many financial crises in EMDEs.

Fourth, depending on the prevalence of their use domestically, stablecoins import the monetary policies of the fiat currencies in the basket that may not be optimal for most EMDEs and could thus impinge on their monetary policies. “Stablecoin-isation” could mean less effective monetary transmission and, in the extreme, countries that face shocks—political, economic or financial—could face deposit outflows from banks and capital flight. This would amplify instability and render policy measures less effective. Countries with

Table 3 Particular challenges of stablecoins for EMDEs

Development challenges	<ul style="list-style-type: none"> • Weaker capacity to address AML/CFT risks • Lack of robust cash-in / cash-out solutions
Macroeconomic challenges	<ul style="list-style-type: none"> • Volatility to local currency • Higher risk of losing monetary control
Cross-border challenges	<ul style="list-style-type: none"> • Capacity constraints in cross-border coordination • Oversight challenges as "host"

large cross-border inflows in stablecoins may face difficulties in maintaining international reserves in hard fiat currencies. This has implications for the functioning of FX and interbank markets, which are shallower in EMDEs. Liquidity and redemption shocks may thus create disruptive spillovers.

Fifth, in light of the different roles discussed earlier, the various entities in stablecoin arrangement are inter-dependent for the overall system to provide smooth and resilient services. Thus, disruptive spillover and spillback effects may emerge. This calls for a consolidated oversight approach to detect and mitigate risks. Such oversight may be impeded by cross-border challenges if entities operate in different jurisdictions. Stablecoin arrangement however may combine elements of multiple regulatory frameworks, e.g. for payment systems, bank deposits, e-money, commodities, FX and securities. In some jurisdictions, there may be gaps as no specific framework would apply. This may create an unlevel playing field if countries adopt different regulatory approaches and impede a holistic regulatory and supervisory approach. EMDEs may have more difficulty to allocate proper resources to adjust their policy frameworks, adopt proportionate supervision and engage in coordination across borders. Moreover, crypto-asset activity currently resides mostly outside the regulatory, supervisory and safety net perimeters. This raises the spectre of domestic regulatory arbitrage and may lead to the build-up of risks related to financial stability (including due to cyber and operational risks), financial integrity and consumer protection which could create confidence spillovers.

Sixth, given reach, scale, network, and “winner takes all” effects, EMDEs will likely act as a “host” to entities in a stablecoin system that provide critical services such as governance and reserve asset management, which may be headquartered elsewhere. Residents in EMDEs may also rely on exchange and custody functions from cross-border VASPs such as exchanges or wallets which may elude “host” supervisory reach. This may call for additional tools for “host” supervisors to regulate cross-border VASPs that offer products or services in their jurisdiction, as the FATF has done in its amended rules. Furthermore, stablecoins have a higher potential from a “host” perspective to become systemically important, even if they are not systemic in a “home” jurisdiction. This could create a misalignment of incentives between “home” and “host” supervisors and impede holistic oversight. This resembles existing challenges posed by supervisory colleges and crisis management groups of financial institutions that are active in small economies. As such, authorities may lack control over the broader stablecoin arrangement and its operations that involve residents. When domestically adopted at scale, this could inhibit monitoring of risks and effective oversight of payments to

prevent illicit use and to foster financial stability, as outlined by international standards. Moreover, it raises questions on consumer protection and redress mechanisms.

Early impressions from interactions with EMDE policy makers yield the following observations around stablecoins:

- The need for an internationally recognised classification and guidelines for legal and regulatory frameworks to identify and address regulatory gaps and the potential for international arbitrage, particularly given that stablecoins could fall under different regulatory classifications.
- The need to review coordination mechanisms to enable a comprehensive and consistent regulatory and supervisory approach across a fragmented ecosystem
- The need for data and information exchange to allow regulators to get a comprehensive view and evaluate whether collaboration arrangements are adequate.

Many of these challenges can be addressed, or at least mitigated, by adequate policies. These could include additional resources on AML/CFT supervision, regulations to limit currency mismatches and further international coordination. Existing frameworks like the Principles for Financial Market Infrastructures (PFMI) can also help address risks (CPMI-IOSCO 2012).¹¹ Moreover, authorities can learn from regulatory and supervisory arrangements of existing financial market infrastructures that operate across borders. For example, the Southern African Development Community (SADC) Payment System Oversight Committee (PSOC) works together to advance the objectives of payment, clearing and settlement systems. In this regard, the SADC PSOC collaborates with various stakeholders to ensuring the safety and efficiency of the regional payment system. As another example, authorities can learn from established frameworks such as the Joint Forum Principles (BCBS 2012) to supervise financial conglomerates that operate across borders and often face regulatory gaps and blind spots. Many of these principles are broadly relevant to stablecoin arrangements, including the need for adequate supervisory powers, supervisory tools that induce timely corrective actions, cross-border coordination mechanisms between supervisors, and corporate governance frameworks. Yet such policies and frameworks take time

¹¹ The PFMI are the international standards for financial market infrastructures, i.e. payment systems that are systemically important, central securities depositories, securities settlement systems, central counterparties and trade repositories.

and resources to be developed and enacted, and the potential opportunities from stablecoins have to be weighed against the substantial risks.

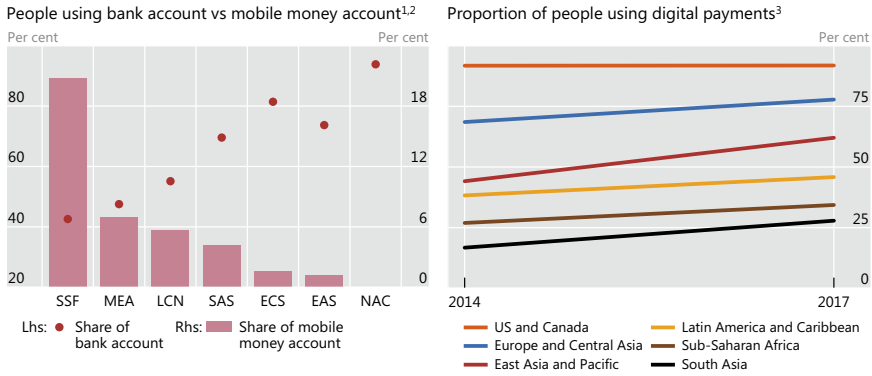
CBDCs—and in particular retail CBDCs—present their own policy challenges for EMDE authorities. In particular, there is a risk that in periods of systemic stress, households and other agents may suddenly shift from bank deposits or other instruments into the CBDC, spurring a “digital run” of unprecedented speed and scale (CPMI/MC 2018: 16). Numerous ideas for capping balances in CBDCs or restricting convertibility between CBDCs and deposits are being proposed (see e.g. Kumhof and Noone 2018; Bindseil 2020). Yet as EMDE authorities can attest, measures to suspend convertibility and restrict retail payment options for the sake of domestic stability are not without their own challenges and drawbacks.

5 Technological Advances Are Already Enhancing Inclusion and Efficiency

Stablecoins and CBDCs are certainly not the only game in town. In recent decades, technological advances have given EMDEs an opportunity to “leapfrog” into the digital economy (IMF and World Bank Group 2018). Fintech facilitates the digitisation of money, making accounts and payments services more accessible, safer, cheaper, more convenient and closer to real time. Across all levels of economic development, the share of unbanked adults and the costs of remittances are falling. Several factors have facilitated these developments.

First, there is a global rise of non-bank e-money issuers such as e-commerce platforms or telecom operators with large user bases that benefit from network effects. E-money is a bridge to commercial bank money, as in most countries it needs to be fully covered by commercial bank money. E-money can be conveniently stored on and exchanged from a mobile phone or online and funds can be transferred through digital channels as well as physical agent locations. This is better suited for many consumers in EMDEs, particularly for those who live in remote areas. In Sub-Saharan Africa, the share of adults with an e-money or mobile money account nearly doubled from 2014 to 2017, to a level of 21% (Graph 8, left panel).¹² Globally, 52% of adults used digital payments in 2017, up from 42% in 2014 (Graph 8, right panel; Demirgüç-Kunt et al. 2018).

¹² Mobile money is used here to mean a form of payment accessible through a mobile device—one form of e-money. Mobile money is sometimes used more broadly to mean the provision of financial services through a mobile device.



¹ EAS = East Asia and Pacific; ECS = Europe and Central Asia; LCN = Latin America and Caribbean; MEA = Middle East and North Africa; NAC = Northern America; SAS = South Asia, SSF = Sub-Saharan Africa. ² Data for 2017. ³ 2017 data for Middle East and North Africa is 38%.
 Source: World Bank Findex data.

Graph 8 Digital payments like mobile money are already taking off (Source World Bank Findex data)

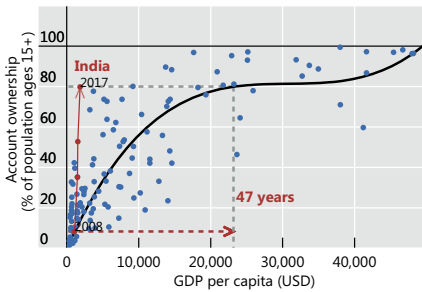
Second, policy makers are facilitating fintech innovation and adoption by updating policy frameworks and promoting digital literacy. Many countries are working on digital ID systems, which provide the opportunity to bring over one billion undocumented people into the financial sector and promote transaction security. The experience with Aadhaar in India is particularly instructive (D’Silva et al. 2019). The combination of digital ID and other services (the “India stack”) has allowed India to lower the cost of KYC checks and increase account ownership from 20% in 2008 to 80% in 2017. One rough estimate, based on cross-country experience, is that it would have taken 47 years to achieve this level of adults with a bank account if India had solely relied on traditional growth processes (Graph 9, left panel).

Third, authorities are upgrading payment infrastructures with “fast payments”, allowing banks and eligible non-banks to offer 24/7, near real-time payments (Bech et al. 2017; 2020). These fast payment systems are now available in over 55 countries (Graph 9, right panel) and show a logistic rate of adoption, similar to the earlier experience with real-time gross settlement (RTGS) systems. Moreover, “open banking” initiatives allow for third party-initiated payment services,¹³ often de-coupling transaction accounts from banks and empowering customers. This can help boost competition.

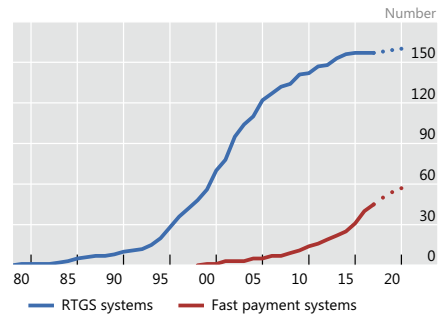
Fourth, feeling the pressure to innovate, incumbent banks and payment providers are embracing fintech to improve their services so consumers can

¹³ Open banking refers to a system in which financial institutions’ data can be shared for users and third-party developers, e.g. through application programming interfaces.

Account ownership rises with income, but countries can leapfrog



Diffusion of fast payments¹



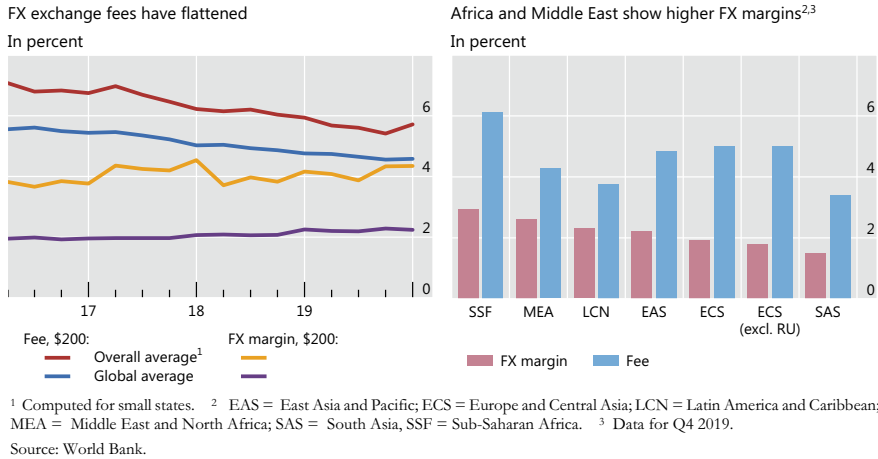
¹ The dashed part of the lines corresponds to projected implementation.

Source: Bech, Shimizu and Wong (2017); FIS (2018); IMF, *World Economic Outlook*, October 2019; World Bank Findex data; Instapay; national data.

Graph 9 Digital technologies can help support inclusion and convenience (Source Bech et al. 2017; FIS 2018; IMF, *World Economic Outlook*, October 2019; World Bank Findex data; Instapay; national data)

conduct payments more conveniently, faster and 24/7. For example, many incumbent banks are joining hands, in some cases also with non-banks, to develop fast payment networks and offer access to their deposit-based products via mobile apps (Petralia et al. 2019). Existing MTOs are increasingly supporting a wide variety of payment instruments and integrating into payment systems in sending and receiving countries—including in some cases with fast payment systems. Central Banks are also increasingly considering extending access to public payment systems to fintech players and operate them on a 24*7 basis.

Finally, new fintechs have extended the MTO model for cross-border transfers by connecting to local payment infrastructures and banks or e-money providers on both sides of a transaction. Closely related to this trend, a range of specialised providers have entered the market establishing non-branded (“white-label”) cross-border payment services (Earthport, MFS Africa and Currencycloud). Incumbent institutions and fintechs can integrate with these white-label solutions to rapidly offer cross-border payment services to their clients. Further, the global financial messaging network SWIFT has launched the Global Payments Initiative (SWIFT gpi) to bring transparency, speed and reliability to correspondent banking transactions. These initiatives could bring down fees in cross-border payments, such as FX fees (Graph 10, left panel). While these fees have come down a bit recently, they remain high for some regions, particularly Africa and the Middle East (Graph 10, right panel).



Graph 10 FX margins make up the bulk of overall fees (Source World Bank)

6 Conclusion

Stablecoin arrangements aspire to improve financial inclusion and cross-border remittances—but they are neither necessary nor sufficient to meet these policy goals. They are not yet tested at scale, and it is unclear whether they would offer lasting competitive advantages over rapidly evolving digital payments services that are built on top of, or aim to improve the existing financial plumbing. Innovations such as digital ID, e-money, mobile banking, open banking and faster payment systems may be adequate in a domestic setting. The development of SWIFT gpi and the cross-border integration of faster payment systems could help improve cross-border payments, although more work is clearly needed.

Meanwhile, stablecoins face various challenges and pose new risks, particularly in EMDEs. Thus authorities may consider to limit or even prohibit the use of stablecoins as a means of payment, and bar regulated entities such as banks and agent networks from holding stablecoins or offering stablecoin services.

Some countries have begun to accelerate their investigations into a CBDC for consumers. However, a new digital equivalent to cash also raises various challenges for EMDE authorities. While research is ongoing, it is not yet clear whether CBDCs are necessary or desirable for all jurisdictions.

Taken together, perhaps the most important contribution of stablecoins thus far is that they have drawn greater—and much-needed—attention to the challenges of financial inclusion and more efficient cross-border payments and remittances. This highlights the efforts underway to strengthen monetary

and financial stability frameworks; promote an enabling regulatory environment for fintech; upgrade payment infrastructures, particularly across borders and ensure a global regulatory level playing field through greater collaboration.

Acknowledgments We would like to thank Raphael Auer, Stijn Claessens, Sebastian Doerr, Matei Dohotaru, Alfonso Garcia Mora, Leonardo Gambacorta, Keith Hart, Marc Hollanders, Yira Mascaro, Leandro Medina, Fritz Schneider, Hyun Song Shin and Mahesh Uttamchandani for helpful comments and suggestions. We thank Haiwei Cao, Giulio Cornelli and Alexandra End for excellent research assistance. This chapter is based on Feyen et al (2020). The views in this chapter are those of the authors only and do not necessarily reflect those of the World Bank Group or the Bank for International Settlements.

References

- Adrian, T., and T. Mancini-Griffoli. 2019. The rise of digital currency. VoxEU, 28 August.
- Arner, D., R. Auer, and J. Frost. 2020. Stablecoins: risks, potential and regulation. Banco de España. *Financial Stability Review*, November.
- Auer, R., and R. Böhme. 2020. The technology of retail CBDCs. *BIS Quarterly Review*, March.
- Auer, R., G. Cornelli, and J. Frost. 2020. Rise of the central bank digital currencies: Drivers, approaches and technologies. BIS Working Paper No 880.
- Bank for International Settlements. 2018. Cryptocurrencies: Looking beyond the hype. BIS Annual Economic Report, June.
- Bank for International Settlements. 2020. Central banks and payments in the digital era. BIS Annual Economic Report, June.
- Basel Committee on Banking Supervision. 2012. Joint forum principles for the supervision of financial conglomerates. September.
- Bech, M., and R. Garratt. 2017. Central bank cryptocurrencies. *BIS Quarterly Review*, September.
- Bech, M., J. Hancock, and W. Zhang. 2020. Fast retail payment systems. *BIS Quarterly Review*, March.
- Bech, M., T. Shirakami, and P. Wong. 2017. The quest for speed in payments. *BIS Quarterly Review*, March.
- Bech, M., U. Faruqui, F. Ougaard, and C. Picillo. 2018. Payments are a-changin' but cash still rules. *BIS Quarterly Review*, March.
- Bech, M., and C. Boar. 2019. Shaping the future of payments, November.
- Bindseil, U. 2020. Tiered CBDC and the financial system. ECB Working Paper Series, No.0 2351, January.

- Boar, C., H. Holden, and A. Wadsworth. 2020. Impending arrival—A sequel to the survey on central bank digital currency. BIS Paper No 107.
- Boar, C., and A. Wehrli. 2021. Ready, steady, go?—Results of the third BIS survey on central bank digital currencies. BIS Paper, forthcoming.
- Brosens, T., and C. Cocuzzo. 2019. Where a retail digital currency could thrive. ING THINK Economic and Financial Analysis, December.
- Brunnermeier, M.K., and D. Niepelt. 2019. On the equivalence of private and public money. *Journal of Monetary Economics* 106: 27–41.
- Carstens, A. 2019a. The future of money and payments, Speech at the Central Bank of Ireland, Whitaker Lecture, March.
- Carstens, A. 2019b. The future of money and the payment system: What role for central banks?, Speech at Princeton University, December.
- Carstens, A. 2021. Digital currencies and the future of the monetary system. Speech at the Hoover Institution, January.
- Committee on Payments and Market Infrastructures. 2016. Correspondent banking, July.
- Committee on Payments and Market Infrastructures. 2019. CPMI quantitative review of correspondent banking data, May.
- Committee on Payments and Market Infrastructures and International Organization of Securities Commissions. 2012. Principles for financial market infrastructures, April.
- Committee on Payments and Market Infrastructures and Markets Committee. 2018. Central bank digital currencies, March.
- Crowe, C., and E. Meade. 2007. The evolution of central bank governance around the world. *Journal of Economic Perspectives* 21 (4): 69–90.
- D’Silva, D., Z. Filkova, F. Packer, and S. Tiwari. 2019. The design of digital financial infrastructure: Lessons from India. BIS Paper No. 106.
- Demirgüç-Kunt, A., L. Klapper, D. Singer, S. Ansar, and J. Hess. 2018. The global fintech database 2017: Measuring financial inclusion and the fintech revolution, April.
- Feyen, E., J. Frost, and H. Natarajan. 2020. Digital money: Implications for emerging market and developing economies. voxEU column, January.
- Financial Action Task Force. 2019a. Money laundering risks from “Stablecoins” and other emerging assets, October.
- Financial Action Task Force. 2019b. Public statement on virtual assets and related providers, June.
- Financial Stability Board. 2017. FSB correspondent banking data report, July.
- Financial Stability Board. 2019. FSB action plan to assess and address the decline in correspondent banking: Progress report, May.
- Financial Stability Board. 2020. Regulation, supervision and oversight of ‘global stablecoin’ arrangements, October.
- Foley, S., J. Karlsen, and T. Putnins. 2019. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies* 32 (5): 1798–1853.

- G7 Working Group on Stablecoins. 2019. Investigating the impact of global stablecoins, October.
- GSMA. 2019. 2018 state of the industry report on mobile money, February.
- Hart, K. 1973. Informal income opportunities and urban employment in Ghana. *The Journal of Modern African Studies* 11 (1): 61–89.
- Hileman, G. 2015. The bitcoin market potential index. Proceedings of the 2nd Workshop on Bitcoin Research in Association with Financial Cryptography and Data Security, Springer.
- International Monetary Fund. 2017. Recent trends in correspondent banking relationships: Further considerations. Policy Paper, April.
- International Monetary Fund. 2020. Philippines: Staff report for the 2019 article iv consultation, February.
- International Monetary Fund and World Bank. 2018. The Bali fintech agenda, October.
- Kahn, C.M., and W. Roberds. 2009. Why pay? An introduction to payments economics. *Journal of Financial Intermediation* 18 (1): 1–23.
- Khiaonrong, T., and D. Humphrey. 2019. Cash use across countries and the demand for central bank digital currency. IMF Working Paper.
- Kumhof, M., and C. Noone. 2018. Central bank digital currencies—Design principles and balance sheet implications. Bank of England Working Paper No. 725.
- Listfield, R., and F. Montes-Negret. 1994. Modernizing payment systems in emerging economies. World Bank Policy Research Paper No 1336.
- Niepelt, D. 2018. Reserves for all? Central bank digital currency, deposits, and their (non)-equivalence. CEPR Discussion Paper 13065.
- Medina, L., and F. Schneider. 2019. Shedding light on the shadow economy: A global database and the interaction with the official one. CESifo Working Paper No 7981.
- Moin, A., E. Siner, and K. Sekniqi. 2019. A classification framework for stablecoin designs. [arXiv:1910.10098v1](https://arxiv.org/abs/1910.10098v1).
- Petralia, K., T. Philippon, T. Rice, and N. Véron. 2019. Banking disrupted? Financial intermediation in an era of transformational technology. Geneva Report 22, September.
- Ratha, D., S. De, E.J. Kim, S. Plaza, G. Seshan, and N.D. Yameogo. 2019. Data release: Remittances to low- and middle-income countries on track to reach \$551 billion in 2019 and \$597 billion by 2021. World Bank People Move blog, October.
- Recorded Future. 2020. How North Korea revolutionized the internet as a tool for rogue regimes, February.
- Rice, T., G. von Peter, and C. Boar. 2020. On the global retreat of correspondent banking. *BIS Quarterly Review*, March.
- Transparency International. 2020. Corruption perceptions index 2019, May.
- World Bank. 2018. The decline in access to correspondent banking services in emerging markets: Trends, impacts and solutions, March.

World Bank. 2019. Remittance prices worldwide, December.

Zetsche, D., D. Arner, and R. Buckley. 2020. Regulating libra: The transformative potential of Facebook's cryptocurrency and possible regulatory responses. *Oxford Journal of Legal Studies*, 36, December.



Blockchain and Other Distributed Ledger Technologies in Finance

Gilles Hilary and Laura Xiaolei Liu

1 Introduction

A blockchain is a form of distributed ledger technology (DLT). Although this technology has grown in prominence, its full potential in terms of practical financial applications and academic research is not completely understood. Blockchain technology was first implemented in 2008 as the underpinning of Bitcoin, which was the first cryptocurrency—a new concept at the time. What started as an esoteric idea to merge elements of finance, game theory, cryptography, and computer science has attracted mainstream attention in the last few years. For example, a study of Google queries from 2008 to 2018 for the word blockchain were rare before February 2013, but their frequency increased gradually until they peaked in December 2017. The pattern in the PRC has been similar, although the number of Baidu searches for the term blockchain peaked a bit later, in March 2018. Naturally, academic interest has followed a similar trend. A review of the number of studies

G. Hilary (✉)

Georgetown University, Washington, DC, USA

e-mail: gilles.hilary@georgetown.edu

L. X. Liu

Guanghua School of Management, Peking University, Beijing, People's Republic of China

e-mail: laura.xiaolei.liu@gsm.pku.edu.cn

of blockchains indexed by Google Scholar every year since 2008 (across all fields). Figure 1 shows that the number has exploded since 2015. Money has also followed a similar trend. According to Coindesk, the increase in investment in blockchain projects slightly predated the popular interest but, the pattern of investments in blockchain ventures also shows a similar pattern.

Although many sectors take a keen interest in blockchains (and other DLTs), at the time that we are writing this chapter, the fall of 2018, supply chain management and financial applications are the most common applications of blockchain technologies. Babich and Hilary (2019a, b) review the application of DLTs to operations (and more generally to enterprise applications) and the research opportunities for academics in that field. They note that fully deployed applications are still rare and many projects fail, but proofs of concepts (PoCs) are becoming increasingly prevalent in the fields of application.

At its core, a blockchain is a new database technology. We elaborate more on its technical details later, but we first consider its strengths and weaknesses. Babich and Hilary (2019b) identify five positive traits of DLTs (see Table 1): *visibility, aggregation, validation, automation, and resiliency*.



Fig. 1 Bitcoin prices

Table 1 Strengths and weaknesses of blockchain technology

Blockchain strengths	Blockchain weaknesses
Visibility	Lack of privacy
Aggregation	Lack of standardization
Validation	Garbage in, garbage out
Automation	Black box effect
Resiliency	Inefficiency

Visibility is the ability of network participants to follow information through a series of transactions. Aggregation ensures that the information in a blockchain can originate from a variety of sources: firms, customers, regulators, and smart sensors. Validation refers to the fact that information captured in a distributed ledger is difficult to tamper with. Automation is the ability to execute certain transactions automatically in response to pre-specified conditions. Resiliency is the feature whereby an entire blockchain database can withstand certain shocks (physical shocks, in particular) better than traditional databases can (see Hilary 2018, 2021 for a review of security issues).

However, Babich and Hilary (2019b) are quick to point out that those advantages are balanced by five corresponding weaknesses: *lack of privacy*, *lack of standardization*, the “*garbage in, garbage out*” (GIGO) problem, the *black box effect*, and *inefficiency*. Lack of standardization stems from the fact that a blockchain is not a unique technology but an umbrella describing a portfolio of protocols that are not yet stable and fully integrated. That lack of standardization can be an issue for financial institutions that for legal reasons must keep their records for decades. Lack of privacy stems from the difficulty in erasing data in blockchain networks; this may make compliance with privacy regulations difficult. The black box effect for blockchains refers to the fact that a blockchain can remove the need to trust a counterparty in some circumstances but at the same time requires “meta-trust” in the blockchain concept. For example, retail investors and regulators need to trust the integrity of the process without necessarily understanding the technical underpinnings. Blockchain inefficiency stems from current limitations in the technology. The GIGO problem stems from the fact that a blockchain is only as good as the information recorded in it. Paradoxically, the difficulty with modifying blockchain data may also make the correction of errors more arduous.

The financial world has applied blockchain technology in at least three main domains (and more applications may surface in the future): (1) the trading of crypto-assets that are standardized and treated like quasi-currencies, (2) the development of mechanisms to raise money by tying the crypto-assets to underlying assets, such as cash-flow rights, and (3) operational functions, such as mechanisms to transfer and store documentation. For example, HSBC Holdings realized the first trade finance transaction using a blockchain by issuing a letter of credit to Cargill in 2018.¹ The development of DLTs may impact operational activities such as custody, compliance

¹ <https://www.cnn.com/2018/05/14/hsbc-makes-worlds-first-trade-finance-transaction-using-blockchain.html>.

(e.g., know your customer [KYC] procedures), or correspondent banking. However, whether and when these back-office tasks can be transferred to a DLT platform remain open questions. For example, in 2018, SWIFT released the results of a PoC in a sandbox environment that allowed the functional transfer of information. However, the report highlighted numerous issues that needed to be solved before the pilot could become fully operational in a real-world setting. Although operational aspects are economically important, we largely ignore them in this chapter because the topic is typically not the focus of academic finance research.

The rest of this chapter proceeds as follows. The second section focuses on the blockchain technology itself and the technical aspects of its first application, Bitcoin. The third section examines the issuance of crypto-assets. We start with cryptocurrencies (and their impact on monetary policy), move to initial coin offerings (ICOs), and finally discuss the issuance of other crypto-assets that are backed by a wide range of underlying assets. The fourth section analyzes the post-issuance behavior of crypto-assets. We explore market integrity, market efficiency, and market microstructure issues. The fifth section considers the token economy. We present our conclusions in the sixth section.

2 What Is a Blockchain?

2.1 The Technology

Blockchain was first introduced in 2008 as the technology supporting Bitcoin, the first successful virtual currency system that eschewed a central authority for issuing currency, transferring ownership, and confirming transactions (see Hilary 2021 for a discussion of the technology). Since then, blockchains have found many applications beyond their initial purpose and many different versions have been implemented. A blockchain is a new form of database technology known as a “distributed ledger.” Traditional centralized databases hold only one master version at any given time. In contrast, a distributed database (as opposed to a distributed ledger) involves multiple nodes (computers) that cooperate under one umbrella to maintain the integrity of the database should one of the nodes fail (for example, experience a hardware problem). However, this approach assumes that no node will alter or manipulate content. That feature is known as “fault tolerance.”

Distributed ledgers (such as blockchains) replicate the database, meaning that each node has an active copy. These versions can temporarily diverge

from one another, but the technology constantly ensures that the different versions converge to a consensus version. A DLT such as a blockchain addresses the possibility that nodes may generate arbitrary data while posing as an honest actor. In other words, they remove the need for nodes to trust each other. However, this feature requires trust in the technology itself and its correct implementation. The goal is for the system to be “Byzantine fault-tolerant,” which is the ability to handle component failure when there is uncertainty about the behavior of any component and its possible failure.

Distributed ledgers can either be decentralized (granting equal rights, within the protocol, to all participants) or centralized (giving certain users particular rights). A blockchain is a specific form of distributed ledger that has a specific data structure, but other forms of distributed ledgers exist. Blockchains can be either public or private. Everyone can join a public blockchain, but participants in a private blockchain must be vetted. Most cryptocurrencies (e.g., Bitcoin) are based on a public blockchain, whereas many enterprise applications rely on private blockchains.²

2.2 The Bitcoin Example

Bitcoin was the first public blockchain (see Harvey 2014, 2016 for an early introduction to the crypto-currency and crypto-finance, and Böhme et al. 2015 for an early economic analysis of Bitcoin). Bitcoin enables a distributed digital ledger to record transactions between two parties in a verifiable and permanent way. Bitcoin’s blockchain is essentially a combination of peer-to-peer software with a public key cryptographic tool. The blockchain contains two kinds of records: transactions and blocks. Blocks hold batches of valid transactions and the additional information necessary for the system to work. To perform this validation, nodes use specific algorithms. These blocks are ordered in a single chain that is replicated over the entire network. Once verified, the transaction is combined with other transactions to create a new time-stamped block of data for the ledger. The new block is then added (by consensus) to the existing blockchain in a way that is meant to be (largely) unalterable. Many of the features present in Bitcoin are common to other distributed ledgers, but most of them (e.g., pseudo-anonymization, traceability, auditability, and immutability) are not intrinsic characteristics of distributed ledgers.

² Blockchain applications adopted and managed by a select group of enterprises are sometime called “consortium blockchains.”

The current consensus mechanism in the Bitcoin blockchain uses the proof-of-work (PoW) approach to stamp blocks. In that approach, the Bitcoin blockchain determines the order of the blocks by using a mathematical “lottery” that is based on solving a mathematical problem (a “cryptographic hash”). The node (or “miner”) that first completes the block can add it to the consensus. However, solving the hash function problem is costly, requiring, among other things, computer time, bandwidth, and electricity. To incentivize miners to incur these costs, the first node that solves the problem (and thus adds a block to the chain) receives a financial reward in the form of newly minted Bitcoins. However, Hinzen et al. (2019) point out some of the economic limitations of the PoW technology.

Various blockchain platforms have made different implementation choices, each with its advantages and drawbacks. For example, an alternative consensus mechanism design may have lower costs but suffer from other technical issues. Public blockchains (and, to a lesser extent, private blockchains) currently face performance and scalability issues. Bitcoin, for example, can handle approximately seven transactions per second (usually fewer), whereas VisaNet can allegedly handle up to 65,000 transactions per second.³ Other DLTs offer the possibility of greater scalability but have not yet been extensively deployed and tested. Chiu and Koepl (2017) and Huberman et al. (2017) find that the current Bitcoin system generates a large social welfare loss, and they explore ways of mitigating that cost.

2.3 Smart Contracts and Tokens

Records in blockchain databases can contain any number of elements, including executable software often called “smart contracts” that can facilitate, verify, or enforce the negotiation or performance of a contract. However, “smart contract” is often a misnomer because self-executable sets of instructions (i.e., lines of codes) will not typically create legally binding contracts.⁴ By allowing the addition of redacted materials to its database, Corda, a blockchain platform commonly used in financial applications, increases the likelihood that a contract integrated on its platform will be legally binding.

Relatedly, the elements contained in the database may lead to “tokenization,” which is the process of converting rights to an asset into a digital token (or record) on a blockchain. Bitcoins and other cryptocurrencies are types of

³ <https://usa.visa.com/dam/VCOM/download/corporate/media/visanet-technology/aboutvisafactsheet.pdf>.

⁴ Smart contracts are called « chain codes » in the Hyperledger environment.

tokens. However, other assets can be similarly tokenized. The tokenization process can, for example, facilitate the trading of illiquid assets and enable micropayments. Tokenization enables an ICO, a financial mechanism that takes advantage of blockchain technology to raise funds.

3 Crypto-Asset Issuance

A DLT facilitates the issuance of different types of assets. In this section, we discuss several cases.

3.1 Cryptocurrency Issuances

As already noted, Bitcoin was the first cryptocurrency. The concept paper was published in 2008, but the first coin was not created until 2009. The first transaction occurred in 2010.⁵ The volume of bitcoin transactions gradually increased until it peaked in late 2017 at more than 400,000 per day; it has since declined to fewer than 250,000 transactions per day (as of the fall of 2018). The price of a bitcoin increased from an initial value close to zero to a historic high of nearly 20,000 USD in late 2017, but by autumn 2018, it had fallen back to less than 4000 USD. We show the historical price of a bitcoin in Fig. 1.⁶

By 2011, other cryptocurrencies had started to appear, using alternative technologies that provided different features. For example, some (e.g., Monero) offer enhanced privacy to execute transactions. By autumn 2018, more than 2000 cryptocurrencies existed. On October 8, 2018, the website coinmarketcap.com indicated an overall market capitalization of the 2042 listed currencies as amounting to 221 billion USD (115 billion for Bitcoin and 23 billion for Ethereum, the second most valuable cryptocurrency). In comparison, the market capitalization of Apple during the same period was approximately 1000 billion USD. Vietnam and Portugal each had GDPs of approximately 220 billion USD in 2017.

The development of cryptocurrencies is creating new policy challenges and opening new areas for research. Most tokens can only be redeemed within their own ecosystem (e.g., FileCoin).⁷ In contrast, cryptocurrencies such as Bitcoin represent a special type of crypto-asset: owners can exchange them

⁵ <https://www.cnn.com/2018/06/18/blockchain-what-is-it-and-how-does-it-work.html>.

⁶ We obtained the Bitcoin price data from Coindesk.

⁷ FileCoin tokens are used to create a storage marketplace. Users pay storage providers for the service with FileCoins. <https://medium.com/swlh/filecoin-and-ipfs-f5e84ac79afa>.

for other products or services outside their own blockchain.⁸ However, even if most crypto-assets do not aim to reach this level of acceptance, they can still be a substitute for fiat money within their own ecosystem. That substitution can be problematic for monetary authorities. For example, Hendry and Zhu (2017) model the coexistence of fiat money and cryptocurrencies and allow the use of both types for different kinds of transactions. Their model shows that the existence of cryptocurrencies can restrict the monetary authorities' ability to control inflation and raises concerns about monetary policy coordination.⁹

Perhaps in response to those concerns, several countries, including the PRC, are considering the possibility of using a digital currency backed by central banks (a central-bank digital currency, or CBDC). Such bank-backed digital currencies are fundamentally different from decentralized cryptocurrencies such as Bitcoin because they benefit from the full support of central banks and are directly controlled by those institutions.¹⁰ A CBDC may play a role in mitigating the zero lower bound issue in interest rate adjustment (Haldane 2015). However, they can create a multiplicity of other problems. For example, Keister and Sanches (2018) analyze a situation in which a CBDC and a private bank's deposits coexist in the marketplace. They find that the introduction of a CBDC often raises general economic welfare but crowds out the banks' deposits and investment. At this time, we do not fully understand the consequences of cryptocurrency development for macroeconomic policy (e.g., optimal monetary policy), if any. These issues offer interesting avenues for research.

3.2 Initial Coin Offerings

As we discussed above, a token is an "object" (or a record) associated with certain rights. Economic agents can exchange tokens to transfer those rights between parties. That process allows for the possibility of ICOs, which are a new method to raise capital. ICOs first appeared in 2014, but there were very few before 2017 (e.g., Giudici et al. 2018). Momtaz (2018a) indicates that the volume of funds raised in ICOs represented 6 billion USD in 2017,

⁸ Several companies, including Overstock, take Bitcoins as payment. In August 2018, a company listed on Austrian Securities Exchange, CCP Technology, issued shares to Singapore-based Penta Global Blockchain Foundation, taking a cryptocurrency, PNT, as payment.

⁹ In a working paper, Lee and Xiao (2018) model how fiscal policy and monetary policy can affect investors' portfolio choices when investors can hold cryptocurrency to circumvent capital control regulations.

¹⁰ The PRC's central bank had filed more than 60 patent applications related to digital currencies as of late 2018.

approximately a fifth of the amount raised in initial public offerings (IPOs) that year. Bourveau et al. (2018) identify 750 ICOs from April 2014 to May 2018 collectively raising 13 billion USD for organizations from 50 countries. Some ICOs are large and skew the distribution. For example, Telegram was able to raise \$1.7 billion, but that amount is exceptional. The vast majority of ICOs are much more modest. For example, Momtaz (2018a) finds that the median amount of money raised in an ICO is less than 6 million USD.

After announcing their decisions to offer an ICO, ventures typically issue a “white paper” describing the technology, the product, the team, and the strategy. The white paper typically states a minimum amount necessary for the ICO to be fully executed, and it may include a maximum amount. Bourveau et al. (2018) mention that in their sample, approximately 85% of ICOs were successful. Investors tender fiat money or crypto-money in exchange for the tokens. Lu (2019) shows that issuing cryptocurrency-denominated ICOs may lead to mispricing as investors improperly associate the fundamental value of the token with that of Bitcoin.

Importantly, the ICOs can be preceded by several rounds of private token issuances. The rights associated with the tokens vary. They can be very similar to stocks, give rights to products (in this case, the tokens are often called “utility tokens”), or, in principle, be any combination of rights that are legally enforceable (e.g., a discount on the price of future shares or products). Giudici et al. (2018) report that only a quarter of ICOs in their sample offered cash flow and governance rights. Interestingly, Giudici et al. (2018) note that 16% of the ICOs offered the opportunity to determine the characteristics of the product/service to be offered (contribution rights). Unlike IPOs, ICO firms do not use an underwriter to help determine the value of the tokens and attract buyers. They may use “advisers”—individuals who publicly endorse the organization but often provide little actual technical advice. Once the process is completed, the tokens are listed on various online platforms (e.g., Bitfinex) and often appear on multiple ones. However, Montaz (2018a) indicates that approximately 13% of ICOs are subsequently delisted from all platforms. He also indicates that the average time from project initiation to the ICO start is 598 days, with a median time of 312 days.

The regulation of these ICOs is currently relatively fluid and varies across jurisdictions. However, the view that ICOs are unregulated is incorrect, at least in major jurisdictions. For example, the United States Securities and Exchange Commission (US SEC) issued a series of pronouncements in 2017 and now holds the view that most ICOs are securities offerings and should

be regulated as such.¹¹ In November 2018, the US SEC announced settled charges against two companies that had sold digital tokens in ICOs. These are the commission's first instances of imposing civil penalties solely for registration violations of ICO securities offerings.¹² Switzerland classifies tokens into three categories: payment (currency) tokens (which are not regulated), asset tokens (which are considered securities), and utility tokens (which are unregulated only if the platform/product behind them is already functional).¹³ As of October 2018, both France and the European Union (EU) were considering a specific form of regulation for ICOs.¹⁴ In 2017, the PRC banned individuals and companies from raising funds through ICOs. Hong Kong allows ICOs but imposes de facto regulation.¹⁵

ICOs offer a new mechanism for raising capital that is distinct from other options, such as going through an IPO on a leading exchange, listing on an over-the-counter (OTC) market, using a crowdfunding platform, receiving funds from a venture capital (VC) funding, and other forms of private financing. However, ICOs are substantially different from VC funding in that there is no bilateral discussion with a supplier of capital; rather, ICOs attempt to raise funds from multiple capital suppliers. For example, Giudici et al. (2018) report an average of 4121 contributors (with a median of 2394) in their sample of ICOs. Chod and Lyandres (2018) compare the use of an ICO with traditional VC financing. They show that ICOs can facilitate risk-sharing between entrepreneurs and investors without transferring control rights, but that such activity may cause underinvestment.

In contrast, IPOs, at least in major jurisdictions, represent a highly regulated activity that often involves large sums of money. Although the situation might change in the future, firms currently engaging in ICOs are very different from those typically engaging in IPOs on a large exchange. As noted above, ICOs (with a few exceptions) involve small ventures. Indeed, many of them do not have revenues or even a finished product to sell.¹⁶ Amsden and Schweizer (2018) cite an Ernst & Young (2017) study reporting that only 5% of ICO ventures had running projects, 11% had prototypes, and 84%

¹¹ <https://hackernoon.com/united-states-ico-regulation-9e2381dee202>.

¹² <https://www.sec.gov/news/press-release/2018-264>.

¹³ <https://cryptovest.com/features/ico-regulations-by-country-a-global-coin-offering-regulatory-overview-may-2018/>.

¹⁴ <https://cointelegraph.com/news/france-finalizes-new-ico-framework-to-attract-innovators-globally>.

¹⁵ <https://cryptovest.com/features/ico-regulations-by-country-a-global-coin-offering-regulatory-overview-may-2018/>.

¹⁶ Signori (2018) notes that 15% of firms launching an IPO in Europe had no revenue prior to the IPO, but that population includes firms listed on secondary markets, such as the Alternative Investment Market (AIM), or on national markets with limited depth, such as Belgium or Denmark.

were merely ideas. However, the behavior of IPOs is well-researched and can constitute a benchmark against which the performance of ICOs can be evaluated. Indeed, Lyandres et al. (2019) point out empirical similarities between IPOs and ICOs. In a recent review of the IPO literature, Lowry et al. (2017) find that the short-term returns in the days following an IPO are consistently positive (close to 20%, on average). The long-term returns are more controversial, with the authors concluding that IPOs underperform the market as a whole but obtain returns comparable to those of similar firms. Studies of over-the-counter markets both in the USA (e.g., Brüggemann et al. 2017; Ang et al. 2013; Eraker and Ready 2015) and in the UK (e.g., Gerakos et al. 2013; Hornock 2015) are consistent with that finding and further suggest that firms do worse when they are listed on illiquid markets than when they are listed on the main markets. The firms have a low probability of graduating to the main market but a high probability of delisting. Returns for most firms are low, although some do spectacularly well.

Overall, ICOs are probably most similar to crowdfunded capital issuances. For example, the US SEC regulation A + , issued in 2015, allows small businesses and startups to raise to \$50 million, which makes this population comparable to the ICO populations. In crowdfunding, entrepreneurs reach out to a potentially large number of investors on selected Internet platforms, such as Kickstarter or Indiegogo. As is the case with ICO investors, crowd-funders can be rewarded with financial compensation, with future products, or with the intrinsic rewards associated with charitable contributions. However, a key difference between crowdfunding and ICOs is that investors in an ICO are free to resell their tokens (upon completion of the lockup period), whereas there may not be a secondary market for financing done through a crowdfunding platform. Furthermore, large ICOs are not fundable through crowdfunding.

The literature on ICOs remains limited. To a large extent, it reproduces the findings of studies on IPOs. For example, preliminary findings suggest that the ICO market behaves like other illiquid and weakly regulated markets. Giudici et al. (2018) report a first-day average return of 930% (with a median return of 25%) for a sample ending in mid-2017. Bourveau et al. (2018) indicate that the median return on the first day is 6% (with a mean of 14%) for a sample ending in mid-2018. Momtaz (2018a) finds similar estimates for the mean (7 to 8%) and a lower number for the median (3%), although he finds an average underpricing of 15% using a larger sample (Momtaz 2018b). Benedetti and Kostovetsky (2018) reach a similar conclusion with larger point estimates (14 to 16%, on average). Long-term returns are always

more difficult to estimate than returns over brief periods. This is particularly true for illiquid markets with a short history. Nevertheless, Bourveau et al. (2018) report a mean return of +39% for ICOs over the first month of trading, but a median of -30%. Momtaz (2018b) reports that although the mean return remains significantly positive over a longer period, investors withholding periods of one month to two years lose approximately 30% from investing in the median firm. Smaller firms outperform larger firms. First-day returns are negatively correlated with long-term returns.

Academic studies have shown that in nearly every market, issuances of capital are more successful when there are more transparency, better governance, higher-quality “guarantors” (e.g., VCs, auditors), and stronger regulations.¹⁷ Giudici et al. (2018), Amsden and Schweizer (2018), Bourveau et al. (2018), and Howell et al. (2019) use multiple proxies to provide evidence for similar patterns in the ICO market. For example, Bourveau et al. (2018) show that weaker information environments are associated with a higher risk of a stock price crash, illiquidity and volatility. Most projects financed by an ICO rely on the existence of a computer code to succeed. The transparency and quality of the code have been consistently shown to be an important aspect of ICO performance. In an interesting paper, Momtaz (2018c) shows that ICOs managed by executives who have a track record of being loyal to their employers fare better than ICOs without such executives. Somewhat relatedly, previous studies show a link between the decision to finance a crowdfunding project and geographic participants (e.g., Agrawal et al. 2015) or social ties (Kuppuswamy and Bayus 2018), but to our knowledge, researchers have not empirically investigated these relations for ICOs.

Indeed, documenting these stylized facts for a new and potentially significant financial market is useful. However, diminishing returns to scale are associated with such documentation. It would be more useful at this point to consider questions that can only be answered in the context of ICOs. For example, most ICOs occur before a product is actually designed, and 16% of tokens give the right to influence the proposed product. Even absent formal rights, the wisdom of the crowd can help to provide important information about future development. Strausz (2016), for example, explains in an analytical study how crowdfunding allows entrepreneurs to contract with consumers before investment. That study describes the tension between increased project screening and moral hazards, but the basic tension could also be analyzed in the context of ICOs. For example, it would be interesting

¹⁷ See Yermack (2017) for a preliminary discussion of the governance issues associated with blockchains.

to analyze the effects of a secondary market and a smart contract on incentive structures. Providing empirical insights on the link between financial and product markets in the context of ICOs would also be relevant. Relatedly, although the vast capital-structure literature has considered the trade-off between equity and debt (with some interest in mezzanine financing), many ICOs provide a new type of instrument on which the cash flows are still contingent by offering rewards directly tied to products but are based on sales rather than on dividends (or earnings).

As discussed above, crowdfunding and ICO platforms offer many similar features but also display certain differences. For example, the seemingly more complex nature of DLTs relative to the more familiar Internet channels could draw or repel certain types of investors. Are there economic benefits and costs associated with one approach versus the other? For example, the open nature of the DLT technology may ensure a greater pool of investors in the secondary market. However, the centralized nature of the crowdfunding platform may give its owners incentives to behave as gatekeepers for removing fraud. In principle, ICOs could offer more liquidity in aftermarket trading, and identifying precisely how this liquidity would behave would be useful.

In addition to providing new research questions about the financing of new ventures, ICO markets offer possibilities for research on other topics. Naturally, that is contingent on their development, which remains somewhat speculative at the time we are writing this chapter. One such area is trade credit.¹⁸ DLTs have the capacity to reduce the friction in supply chains (see Babich and Hilary 2019b). For example, smart contracts coupled with sensors can ensure prompt and guaranteed payments as soon as goods reach their destination. More traceable products and cash flows might allow firms to collateralize their accounts receivable more effectively and to receive payments from banks even faster—for example, through factoring. This faster collateralization can help firms reduce the need for working capital and can lessen the need to conserve cash on the balance sheet to buffer against unexpected delays in collection, which may also affect other dimensions of corporate behavior. For example, if firms do not have to worry as much about the risk of bankruptcy induced by working capital fragility, they may increase other aspects of corporate risk-taking (e.g., innovation).

A second potentially important area is the development of illiquid asset markets. The analytical literature (e.g., Ang et al. 2014; Vayanos and Wang 2012) has noted the importance of liquidity for asset-pricing in a portfolio context, but the ability to “liquify” extends beyond thinly traded stocks. For

¹⁸ Gelsomino et al. (2016) provide a review of the literature on supply chain finance.

example, new classes of assets (such as art, fine wines, and pollution quotas) could become mainstream investments if more effective trading mechanisms become available. As we note above, illiquidity and crash risks plague OTC markets. It would be interesting to analyze how the acceptance of blockchains by exchanges (e.g., the NASDAQ) could affect liquidity, trade efficiency, and ultimately price behavior. Another area of interest for analysis would be the development of internal markets. Distributed ledger technologies can provide a structure for efficiently trading thinly traded assets. One can imagine the trading of claims on internal projects (essentially internal ICOs) or the development of a trading platform for internal resources such as machine hours. Those sorts of trading markets would lead to more efficient use of internal resources.

3.3 Other Crypto-Assets

In addition to cryptocurrencies (which are typically not backed by any other assets) and ICO tokens (which are backed by a specific venture), other assets can now be traded through tokens. Indeed, anything of value can be associated with a token and traded on an exchange that is more or less organized. Examples include gold (e.g., HelloGold), oil (e.g., OilCoin), or even DNA (e.g., Nebula Genomics). In the midst of a devastating financial and economic crisis, Venezuela launched a cryptocurrency (El Petro) backed by national reserves of natural resources. The typical justification for creating these derivative assets is to lower the transaction costs associated with the trading of the underlying assets.

Recently, cryptocurrencies have been created that are backed by fiat money. The goal of these “stable coins” is to offer a relatively stable exchange rate to make the cryptocurrencies an electronic “unit of account” while preserving features of crypto-assets such as the possibility of executing smart contracts. As we noted above, the value of cryptocurrencies such as Bitcoin is very volatile. Tether, the most popular stable coin to date, claims to maintain 100% USD reserves as collateral to guarantee a one-to-one exchange rate pegged to the US dollar. Other stable currencies use alternative mechanisms to maintain the exchange rate. For example, they can over-collateralize other crypto-assets (e.g., DAI) or use another floating-rate cryptocurrency to adjust the supply of the stable coin (e.g., BaseCoin). Recently, the New York Department of Financial Services approved the issuance of two stable coins pegged to the US dollar (GUSD and PAX). The two companies backing those coins

(Gemini Trust and Paxos) obtained the endorsement after convincing the regulator that they would comply with stringent compliance requirements.¹⁹

Recent blockchain projects, such as Valid and GXChain, allow individuals to own and use the data they have generated. Individuals can then obtain compensation from companies that use their data. New technology, such as zero-knowledge proof combined with a blockchain, enables data sharing in a decentralized environment without revealing proprietary information.²⁰ That type of data sharing may lead to the distribution of a large amount of data and the creation of a new asset class. Recent papers have started to investigate the economic consequences of data ownership and sharing. For example, Jones and Tonetti (2018) build a macroeconomic model showing that the allocation of data property rights to consumers may provide a more efficient resource allocation. Easley et al. (2018) analyze the economic consequences of selling consumer data to oligopoly producers. They consider e-commerce settings with three players: a platform (e.g., e-Bay or Taobao), firms (companies that sell their products through the platform), and retail consumers. Their model shows that data ownership by either consumers or platforms can maximize social surplus, whereas firm ownership cannot because competition among firms will lead to suboptimal levels of data sharing.

Research on such alternative crypto-assets is limited (the modeling of the stable coin mechanism by Routledge and Zetlin-Jones, 2018, is an exception) but may be fruitful. For example, understanding how data assets are measured, valued, and shared would be interesting. More work will also be required if we are to understand the optimal design of a stable coin and the potential risks associated with different models.

4 Post-issuance Behavior

4.1 Market Integrity

Historically, unlike ICOs, cryptocurrency exchanges, and especially the coin-to-coin exchanges that do not take fiat money, have been largely unregulated. Criminals have used this new medium of transaction for illicit trading (e.g., Foley et al. 2018) and other criminal enterprises such as money laundering.²¹

¹⁹ For example, Griffin and Shams (2018) provide results suggesting that Techer's USD collateral may be less than 100% and find evidence consistent with this hypothesis. *Official endorsement may mitigate the risk of a similar problem for GUSD and PAX.*

²⁰ See Hilary (2018, 2021) for a discussion of the zero-knowledge proof approach.

²¹ <https://bitcoin.fr/ce-que-dit-tracfin-a-propos-de-bitcoin/>.

However, that lack of regulation may be changing. China banned all crypto-exchanges in 2017. Japan had 16 licensed exchanges by autumn 2018, but fewer than 10 cryptocurrencies are allowed to be traded on those licensed exchanges. In the USA, the CEO of NASDAQ said in April 2018 that it would consider becoming a crypto-exchange in the future.²²

Theft, fraud, and price manipulation affect the markets that remain unregulated, reducing their liquidity. For example, Coinmarketcap lists 225 token exchanges, but only 125 of them had a meaningful trading volume in 2018.²³ Li and Wang (2018) explore pump-and-dump schemes in several cryptocurrency exchanges. They demonstrate the benefits, in terms of liquidity and prices, of fighting those schemes, using the example of Bittrex, a trading platform.²⁴ One of the most famous theft cases was that of Mt. Gox, a platform that was the largest cryptocurrency exchange at the time (2014). After 450 million USD worth of bitcoins disappeared, Mt. Gox shut down. Other platforms (Cointrader, Bitcurex, and Yobit, to list a few) also shut down after falling prey to alleged cyber-hacks.²⁵ Furthermore, theft is not the only market integrity issue. For example, another large exchange, Binance, canceled suspicious trades following a cyber-hack of the platform.²⁶ One of the largest token exchanges, OkEx, canceled a series of Bitcoin future trades that followed a coordinated market manipulation.²⁷ An analysis of transactions from Mt. Gox, Gandal et al. (2018) finds that the suspicious trading activity of a single actor caused massive spikes in Bitcoin's price and trading volume. Griffin and Shams (2018) investigate the trading activity of the stable coin Tether. They show that large Bitcoin purchases with Tether follow market downturns and precede subsequent increases in Bitcoin prices. They report that 50% of the return from Bitcoin and 64% from other top cryptocurrencies are concentrated in 1% of the periods with heavy Tether transactions.

The academic literature has thoroughly studied the market structure of established markets (e.g., Easley and O'Hara 1995) and price manipulation in

²² <https://www.cnbc.com/2018/04/25/nasdaq-is-open-to-becoming-cryptocurrency-exchange-ceo-says.html>.

²³ <https://coinmarketcap.com/exchanges/volume/24-hour/>. We define a meaningful volume above one million USD over the previous 24 hours (website was checked on December 2, 2018).

²⁴ Anecdotal evidence suggests that some exchanges may be inflating their trading volume (e.g., Ribes 2018).

²⁵ One exception is Bitfinex, which claimed to have lost approximately \$69 million of Bitcoin, but later managed to repay its customers.

²⁶ <https://cointelegraph.com/news/binance-reverses-irregular-trades-resumes-trading-amidst-community-confusion-about-hack>.

²⁷ <https://support.okex.com/hc/en-us/articles/360002320751-Details-of-the-Futures-Trading-Incident-on-Mar-30>.

equity and derivative markets (e.g., Aggarwal and Wu 2006; Ni et al. 2005). In contrast, the literature on crypto-markets is more limited. Additional rigorous academic analyses of these crypto-exchanges are warranted, not only for their academic merits but also for their implications for regulators.

4.2 Market Efficiency

Several studies have investigated the pricing of cryptocurrencies. Most of the early empirical studies focused on Bitcoin. Perhaps unsurprisingly, many researchers (e.g., Urquhart 2017; Bariviera 2017; Nadarajah and Chu 2017; Brauneis and Mestel 2018; Detzel et al. 2018) have found multiple examples of inefficiencies in market pricing. Makarov and Schoar (2019) show there exists potential cross-crypto-exchange arbitrage opportunities. Athey et al. (2016) find mixed evidence on whether market fundamentals can explain Bitcoin prices. In particular, they show that as of mid-2015, active use of Bitcoin as a means of payment was not growing quickly, and investors and infrequent users held the majority of Bitcoins. Ciaian et al. (2016) and Liu and Tsyvinski (2018) show that macroeconomic factors do not explain Bitcoin prices well. Gandal and Halaburda (2014) and Hu et al. (2018) show that Bitcoin returns move negatively with other cryptocurrencies. Yu and Zhang (2017) show that the Bitcoin price often violates the law of one price, and they link this finding to capital controls in different countries. Mai et al. (2018) document that posts on social media platforms significantly impact on Bitcoin returns. Several papers, including one by Bukovina and Martiček (2016), find that sentiment affects Bitcoin volatility, and they conclude that Bitcoin is speculative. Yang (2018) tests several price-related anomalies and finds that momentum exists in cryptocurrency markets.

However, one challenge for this line of research is that the baseline model of market equilibrium has not yet been fully characterized, and progress in that direction would be particularly useful. Liu et al. (2019) represent a step in this direction. They show that three factors (i.e., cryptocurrency market, size, and momentum) capture the cross-sectional expected cryptocurrency returns. Given the intensity of speculation in crypto-markets, they also represent an ideal setting for testing predictions stemming from bubble theory. For example, Scheinkman and Xiong (2003) and Hong et al. (2006), among others, propose a “resale option” theory that may be applicable in this setting. Cong et al. (2018) provide a stochastic, rational bubble model of Bitcoin in a macroeconomic framework, and they show that regulation risk and sentiment both affect the value of Bitcoin—a finding that is consistent with their predictions.

4.3 Market Microstructure

Several papers investigate mining protocol designs, largely in the context of Bitcoin. Using different game theory approaches, Eyal and Sirer (2014), Biais et al. (2019), and Cong et al. (2018) consider the mechanisms that lead to token production concentration and collusion in cryptomarkets. The general intuition is that risk-sharing benefits attract independent producers to mining pools and decrease mining diversification. However, cross-pool diversification and endogenous pool fees can mitigate that gradual pool concentration. Easley et al. (2019) also investigate the role of transaction fees in the Bitcoin blockchain's evolution from a mining-based structure to a market-based ecosystem. They also find that microstructure features, such as exogenous structural constraints, influence the dynamics of a crypto-asset environment.

5 Token Economy

5.1 Tokens and Industrial Organization

Distributed ledger technologies have enabled two key innovations: decentralized consensus and tokens. These innovations have important implications for research on industrial organizations. For example, Abadi and Brunnermeier (2018) point out that no ledger can simultaneously satisfy correctness, decentralization, and cost efficiency, although they are the ideal qualities of any record-keeping system. Still, they conclude that a blockchain is more efficient than a traditional ledger managed by a centralized intermediary. Cong and He (2019) show that decentralization through a blockchain facilitates the entry of new competitors but may foster collusion among incumbents. Babich and Hilary (2019b) describe how the concept of *diagonality*, a complex hybrid market structure in which *horizontalities* (e.g., spot markets) and *verticalities* (e.g., supply chain networks) come together, can affect monopoly power.

Several other studies have investigated the importance of native tokens. Cong et al. (2018) focus on the network effect and show the existence of dynamic feedback between user adoption and the responsiveness of token prices to expectations about the future growth of the platform.²⁸ Li and Mann (2018) explore the use of native tokens to overcome coordination

²⁸ Sockin and Xiong (2018) also model a network effect in a blockchain system. They show that two equilibria exist, with either a small or large number of people joining the platform.

failure when a platform needs a critical mass. Catalini and Gans (2018) point out that distributing native tokens can assist entrepreneurs in discovering the consumer valuation of their service.

With the evolution of the industry and the adoption of more applications, additional theoretical work focusing on understanding the economic fundamentals of industrial decentralization would be useful. Examples of interesting questions that link cryptomarkets and industrial organizations include the following. What are the implications of decentralization? Why would a blockchain ecosystem need a native token? What is the optimal token design for the success of a blockchain ecosystem?

5.2 Token Economy and Theory of the Firm

Finally, an understanding of decentralized blockchain communities or ecosystems as forms of organization is potentially very important. For example, the academic literature approaches many corporate finance questions, such as the characterization of agency issues, assuming that firms should maximize the value of their equity and debt. However, in a decentralized blockchain ecosystem, there is no “firm,” even though blockchain participants are similar to corporate stakeholders.²⁹

Magill et al. (2015) explore the stakeholder view of the firm. They show that once property rights have been granted to employees and consumers, boards can instruct managers to maximize the total value for all stakeholders. They argue that “if the firm can issue consumer and worker rights and if these rights can be traded on reasonably liquid markets, then their market prices will reveal the benefits that consumers and workers derive from being stakeholders of the firm” (Magill et al. 2015: 1689). Although Magill et al. do not make explicit references to DLTs, such technologies would enable the implementation of their models. For example, DLTs facilitate the allocation of property rights through token issuances and subsequent trading on platforms that can provide liquid markets for those rights.

Various organizations have already adopted this approach. For example, FCoin, a cryptocurrency exchange, implemented a “trading as mining” mechanism in 2018 that granted consumers (in this case, traders) rights similar to those of equity holders. Traders received a cryptocurrency, FCoin Token (FT for short), based on their trading volume and the most recent FT market price. FCoin distributed most of the transaction fees it charged to FT coin

²⁹ Liu (2018) argues that a blockchain community is a new organizational form in which stakeholders are also the owners of the organization.

holders. The trading platform was initially very successful, but it stopped using this distribution mechanism when the FT price dropped significantly.³⁰ Qutoutiao, a reading application, adopts a “reading as mining” mechanism and allocates tokens to people who use the application to read. The firm distributes a certain percentage of the profits to the token holders.

Practitioners are still pondering whether a traditional corporation can evolve into a decentralized ecosystem, and this is a fruitful line of research for academics. Potential research questions might include the following. How would these decentralized organizations differ from traditional value-maximizing corporations or mutual companies? What would be the impact of decentralized organizations on social welfare and resource allocation? If these decentralized ecosystems become common, they will profoundly impact the theory of the firm.

6 Conclusions

Bitcoin, and the blockchain technology underlying it, have drawn significant attention in recent years. Other crypto-assets and additional distributed ledger technologies have emerged. Practitioners are exploring their applications, but most projects are still at the POC stage, and many have failed. Regulators are looking for the optimal way to regulate these new activities, especially activities related to cryptocurrency issuances and exchanges. The challenges for the regulators are multidimensional and include issues in accounting, taxation, and security. For example, firms typically record cryptocurrencies as intangible assets on their balance sheet. As the cost method is used, companies can cherry-pick realized gains or losses on their income statements. However, dubious quality prices mean that these assets are unlikely to be marked-to-market anytime soon.

In this survey, we first introduce the distributed ledger technologies that enable crypto-assets. We then review questions and findings related to the issuance of crypto-assets (e.g., through ICOs). We highlight research questions in corporate finance regarding how to take advantage of these new settings. In the asset-pricing field, we review previous studies that have focused on market integrity and market efficiency and highlight salient unanswered questions.

³⁰ <https://www.coindesk.com/new-crypto-exchange-draws-fire-over-controversial-business-model>.

However, crypto-assets introduce a new economy in which information can be directly distributed and traded. These assets raise novel and fundamental research questions. For example, the development of various forms of decentralized organization may materially impact the theory of the firm. These DLT-based communities are different from value-maximizing firms in many respects. How they might optimally finance, invest, and govern themselves are open questions. For example, new agency conflicts between the community funders, participants, and other stakeholders may create new challenges.

As new industries emerge and old ones reinvent themselves, academics can play a role in shaping our understanding of the emerging challenges. We hope this work will have a positive impact both on the firms involved in these new developments and on society at large. We expect that academics and practitioners alike will be interested in reading more on these various issues.

Acknowledgments We thank Thomas Bourveau, Will Lin Cong, Allen Huang, Xuewen Liu, Paolo Petacchi, Qiaolin Xiao, Liyan Yang, and Yang Yang for their helpful comments and suggestions. All errors are our own. Liu expresses thanks to the financial support from the Blockchain Lab and Guanghua-Duxiaoman Fintech Lab at the Guanghua School of Management, Peking University, and the Guanghua Thought-Leadership Platform for their financial support.

References

- Abadi, J., and M. Brunnermeier. 2018. Blockchain economics. https://scholarprinceton.edu/sites/default/files/markus/files/blockchain_paper_v3g.pdf.
- Aggarwal, R.K., and G. Wu. 2006. Stock market manipulations. *The Journal of Business* 79 (4): 1915–1954.
- Agrawal, A., C. Catalini, and A. Goldfarb. 2015. Crowdfunding: Geography, social networks, and the timing of investment decisions. *Journal of Economics & Management Strategy* 24 (2): 253–274.
- Amsden, R., and D. Schweizer. 2018. Are Blockchain crowdsales the new ‘gold rush’? Success determinants of initial coin offerings. Available at SSRN: <https://ssrn.com/abstract=3163849>.
- Ang, A., D. Papanikolaou, and M.M. Westerfield. 2014. Portfolio choice with illiquid assets. *Management Science* 60 (11): 2737–2761.
- Ang, A., A.A. Shtauber, and P.C. Tetlock. 2013. Asset pricing in the dark: The cross-section of OTC stocks. *The Review of Financial Studies* 26 (12): 2985–3028.
- Athey, S., I. Parashkevov, V. Sarukkai, and J. Xia. 2016. Bitcoin pricing, adoption, and usage: Theory and evidence. Working Paper. <https://www.gsb.stanford.edu/faculty-research/working-papers/bitcoin-pricing-adoption-usage-theory-evidence>.

- Babich, V., and G. Hilary. 2019a. Blockchain and other distributed ledger technologies in operations. *Foundations and Trends in Technology, Information and Operations Management*. Available at SSRN 3232977.
- Babich, V., and G. Hilary. 2019b. Distributed ledgers and operations: What operations management researchers should know about Blockchain technology. *Manufacturing & Service Operations Management*.
- Bariviera, A.F. 2017. The inefficiency of Bitcoin revisited: A dynamic approach. *Economics Letters* 161: 1–4.
- Benedetti, Hugo E., and Leonard Kostovetsky. 2018. Digital Tulips? Returns to investors in initial coin offerings. Available at SSRN: <https://ssrn.com/abstract=3182169> or <http://dx.doi.org/10.2139/ssrn.3182169>.
- Biais, B., C. Bisiere, M. Bouvard, and C. Casamatta. 2019. The Blockchain Folk Theorem. *Review of Financial Studies* 32 (5): 1662–1715.
- Böhme, R., N. Christin, B. Edelman, and T. Moore. 2015. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives* 29 (2): 213–238.
- Bourveau, T., E.T. De George, A. Ellahie, and D. Macciocchi. 2018. Initial coin offerings: Early evidence on the role of disclosure in the unregulated crypto market. SSRN Working Paper. <https://ssrn.com/abstract=3193392>.
- Brauneis, A., and R. Mestel. 2018. Price discovery of cryptocurrencies: Bitcoin and beyond. *Economics Letters* 165: 58–61.
- Brüggemann, U., A. Kaul, C. Leuz, and I. Werner. 2017. The twilight zone: OTC regulatory regimes and market quality. *Review of Financial Studies* 31 (3): 898–942.
- Bukovina, J., and M. Martíček. 2016. Sentiment and Bitcoin volatility. https://ideas.repec.org/p/men/wpaper/58_2016.html.
- Catalini, C., and J.S. Gans. 2018. Initial coin offerings and the value of crypto tokens. National Bureau of Economic Research, NBER Working Paper No. 24418.
- Chiu, J., and T.V. Koepl. 2017. The economics of cryptocurrencies—Bitcoin and beyond. Working Paper. <https://ideas.repec.org/p/qed/wpaper/1389.html>.
- Chod, J., and E. Lyandres. 2018. A theory of ICOs: Diversification, agency, and information asymmetry. Working Paper. <http://faculty.chicagobooth.edu/workshops/omscience/pdf/Spring%202018/Chod%20Lyandres.pdf>.
- Ciaian, P., M. Rajcaniova, and D.A. Kancs. 2016. The economics of Bitcoin price formation. *Applied Economics* 48 (19): 1799–1815.
- Cong, L.W., and Z. He. 2019. Blockchain disruption and smart contracts. *Review of Financial Studies* 32 (5): 1754–1797.
- Cong, L.W., Z. He, and J. Li. 2018. Decentralized mining in centralized pools. Available at SSRN 3143724.
- Cong, L.W., Y. Li, and N. Wang. 2018. Tokenomics: Dynamic adoption and valuation. Working Paper. <http://files.sitebuilder.name.com/3b/6d/3b6d9eaf-35e8-4e34-8ff8-48f4c921328b.pdf>.

- Detzel, A. L., H. Liu, J. Strauss, G. Zhou, and Y. Zhu. 2018. Bitcoin: Learning, predictability and profitability via technical analysis. Working Paper. http://apps.olin.wustl.edu/faculty/liuh/Papers/Bitcoin_DLSZZ11.pdf.
- Easley, D., and M. O'Hara. 1995. Market microstructure. In *Handbook of finance, handbooks in operations research and management science series*, ed. R. A. Jarrow, V. Maksimovic, and W. T. Ziemba, 9: 375–383. North Holland Press.
- Easley, D., M. O'Hara, and S. Basu. 2019. From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*.
- Easley, D., S. Huang, L. Yang, and Z. Zhong. 2018. The economics of data. Working Paper. <http://www.cafr-sif.com/2018/2018selected/The%20Economics%20of%20Data.pdf>.
- Eraker, B., and M. Ready. 2015. Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks. *Journal of Financial Economics* 115 (3): 486–504.
- Ernst & Young. 2017. EY research: Initial coin offerings (ICOs). Available at [http://www.ey.com/Publication/vwLUAssets/ey-research-initial-coin-offerings-icos/\\$File/ey-research-initial-coin-offerings-icos.pdf](http://www.ey.com/Publication/vwLUAssets/ey-research-initial-coin-offerings-icos/$File/ey-research-initial-coin-offerings-icos.pdf).
- Eyal, I., and E.G. Sirer. 2014. How to disincentivize large Bitcoin mining pools. Blog post. <http://hackingdistributed.com/2014/06/18/how-to-disincentivize-large-bitcoin-mining-pools>.
- Foley, S., J. Karlsen, and T.J. Putniņš. 2018. Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies? SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3102645.
- Gandal, N., and H. Halaburda. 2014. Competition in the cryptocurrency market. Bank of Canada Working Paper 2014-33. Available at <https://www.bankofcanada.ca/wp-content/uploads/2014/08/wp2014-33.pdf>.
- Gandal, N., J.T. Hamrick, T. Moore, and T. Oberman. 2018. Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics* 95: 86–96.
- Gelsomino, L.M., R. Mangiaracina, A. Perego, and A. Tumino. 2016. Supply chain finance: A literature review. *International Journal of Physical Distribution & Logistics Management* 46 (4): 348–366.
- Gerakos, J., M. Lang, and M. Maffett. 2013. Post-listing performance and private sector regulation: The experience of London's alternative investment market. *Journal of Accounting and Economics* 56 (2–3): 189–215.
- Griffin, J. M., and A. Shams. 2018. Is Bitcoin really un-tethered. Available at SSRN. <https://ssrn.com/abstract=3195066> or <http://dx.doi.org/10.2139/ssrn.3195066>.
- Giudici, G., S. Adhami, and S. Martinazzi. 2018. Why do businesses go crypto? An empirical analysis of initial coin offerings. *Journal of Economics and Business* 100: 64–75.
- Haldane, A. 2015. How low can you go. Speech delivered at the Portadown Chamber of Commerce, Northern Ireland. Available at www.mondovisione.com/.../how-low-can-you-go-speech-by-andrew-g-haldane-chief.
- Harvey, C.R. 2014. Bitcoin myths and facts. SSRN Working Paper. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2479670.

- Harvey, C.R. 2016. Cryptofinance. SSRN Working Paper. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2438299.
- Hendry, S., and Y. Zhu. 2017. A framework for analyzing monetary policy in an economy with emoney. Bank of Canada. Available at https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=mmmspr2017&paper_id=104.
- Hilary, G. 2018. Blockchain: Security and confidentiality. *French Military Police Journal* (Revue de la Gendarmerie Nationale) 4 (4): 98–103.
- Hilary, G. 2021. Blockchain and other distributed ledger technologies: An advanced primer. In *Innovative technology at the interface of finance and operations*, ed. V. Babich, J. Birge, and G. Hilary. Springer Series in Supply Chain Management, Springer Nature.
- Hinzen, F.J., K. John, and F. Saleh. 2019. Proof-of-work's limited adoption problem. NYU Stern Working Paper.
- Hong, H., J. Scheinkman, and W. Xiong. 2006. Asset float and speculative bubbles. *Journal of Finance* 61 (3): 1073–1117.
- Hornok, J.R. 2015. The alternative investment market: Helping small enterprises grow public. *Ohio State Entrepreneurial Business Law Journal* 9 (2): 323–377.
- Howell, S.T., M. Niessner, and D. Yermack. 2019. Initial coin offerings: Financing growth with cryptocurrency token sales. *Review of Financial Studies*.
- Hu, A., C.A. Parlour, and U. Rajan. 2018. Cryptocurrencies: Stylized facts on a new investible instrument. Available at SSRN: <https://ssrn.com/abstract=3182113> or <http://dx.doi.org/10.2139/ssrn.3182113>.
- Huberman, G., J.D. Leshno, and C.C. Moallemi. 2017. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. Center for Economic Policy Research, CEPR Discussion Paper No. 12322.
- Jones, C., and C. Tonetti. 2018. Nonrivalry and the economics of data. In 2018 Meeting Papers (No. 477). Society for Economic Dynamics.
- Keister, T., and D. Sanches. 2018. Managing aggregate liquidity: The role of a central bank digital currency. An updated version of this paper is titled Should Central Banks Issue Digital Currency? Available at www.toddkeister.net/pdf/KS_CBDC.pdf.
- Kuppuswamy, V., and B.L. Bayus. 2018. Crowdfunding creative ideas: The dynamics of project backers. In *The economics of crowdfunding: Startups, portals and investor behavior*, ed. D. Cumming and L. Hornuf, 151–182. Palgrave Macmillan.
- Lee, A., and S.X. Xiao. 2018. Digital currency, international portfolio choices and capital control. Working Paper.
- Li, T., and B.L. Wang. 2018. Cryptocurrency pump-and-dump schemes. SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3267041.
- Li, J., and W. Mann. 2018. Initial coin offering and platform building. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2018-af-conference/paper-li.pdf.
- Liu, X. 2018. Blockchain community—A new organization form. *Frontiers*, June 11–17.

- Liu, Y., and A. Tsyvinski. 2018. Risks and returns of cryptocurrency. NBER Working Paper. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3226952.
- Liu, Y., A. Tsyvinski, and X. Wu. 2019. Common risk factors in cryptocurrency. NBER Working Paper. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3379131.
- Lowry, M., R. Michaely, and E. Volkova. 2017. Initial public offerings: A synthesis of the literature and directions for future research. *Foundations and Trends in Finance* 11 (3–4): 154–320.
- Lu, F. 2019. The numeraire effect in initial coin offerings. MIT Working Paper.
- Lyandres, E., B. Palazzo, and D. Rabetti. 2019. Do tokens behave like securities? An anatomy of initial coin offerings. SSRN Working Paper. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3287583.
- Magill, M., M. Quinzii, and J.C. Rochet. 2015. A theory of the stakeholder corporation. *Econometrica* 83 (5): 1685–1725.
- Mai, F., Z. Shan, Q. Bai, X. Wang, and R.H. Chiang. 2018. How does social media impact Bitcoin value? A test of the silent majority hypothesis. *Journal of Management Information Systems* 35 (1): 19–52.
- Makarov, I., and A. Schoar. 2019. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*.
- Momtaaz, P.P. 2018a. Initial coin offerings. Available at SSRN: <https://ssrn.com/abstract=3166709> or <http://dx.doi.org/10.2139/ssrn.3166709>.
- Momtaaz, P.P. 2018b. The pricing and performance of initial coin offerings. Available at SSRN: <https://ssrn.com/abstract=3169682> or <http://dx.doi.org/10.2139/ssrn.3169682>.
- Momtaaz, P.P. 2018c. Initial coin offerings, asymmetric information, and loyal CEOs. Available at SSRN: <https://ssrn.com/abstract=3167061> or <http://dx.doi.org/10.2139/ssrn.3167061>.
- Ni, S.X., N.D. Pearson, and A.M. Poteshman. 2005. Stock price clustering on option expiration dates. *Journal of Financial Economics* 78 (1): 49–87.
- Nadarajah, S., and J. Chu. 2017. On the inefficiency of Bitcoin. *Economics Letters* 150: 6–9.
- Ribes, S. 2018. Chasing fake volume: A crypto-plague. <https://medium.com/@sylvainartplayribes/chasing-fake-volume-a-crypto-plague-ea1a3c1e0b5e>.
- Routledge, B., and A. Zetlin-Jones. 2018. Currency stability using blockchain technology. In Society for Economic Dynamics 2018 Meeting Papers.
- Scheinkman, J.A., and W. Xiong. 2003. Overconfidence and speculative bubbles. *Journal of Political Economy* 111 (6): 1183–1220.
- Signori, A. 2018. Zero-revenue IPOs. *International Review of Financial Analysis* 57: 106–121.
- Sockin, M., and W. Xiong. 2018. A model of cryptocurrencies. Working Paper. <https://www.princeton.edu/~wxiong/papers/Crypto.pdf>.

- Strausz, R. 2016. Mechanism design with partially verifiable information. Cowles Foundation Discussion Paper No. 2040. <http://cowles.yale.edu/sites/default/files/pub/d20/d2040.pdf>.
- Urquhart, A. 2017. Price clustering in Bitcoin. *Economics Letters* 159: 145–148.
- Vayanos, D., and J. Wang. 2012. Theories of liquidity. *Foundations and Trends in Finance* 6 (4): 221–317.
- Yang, H. 2018. Behavioral anomalies in cryptocurrency markets. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3174421.
- Yermack, D. 2017. Corporate governance and blockchains. *Review of Finance* 21 (1): 7–31.
- Yu, G., and J. Zhang. 2017. A revisit of capital controls policies when Bitcoin is in town. SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3053474.



Consensus Mechanisms for the Blockchain

Joshua S. Gans and Neil Gandal

1 Introduction

Economic systems have demand for reliability in terms of information regarding what has happened in the past. The most influential statement of this is from Kocherlakota (1998) who argued that trade beyond two agents requires a record of past performance in trade and showed that money can provide a means of carrying this information in a reliable way. However, the same notions also drive the economic literature on trust and reputation and the economic theory of contracts (Hart 1995). In this regard, blockchain technologies—as they have evolved over the past few decades—represent a new technology that can provide reliable information regarding previous events. As Catalini and Gans (2019) argue, in this respect, from an economic point of view the blockchain can be seen as a potentially large reduction in the cost of verification of past performance.

J. S. Gans (✉)

Rotman School of Management, University of Toronto, Toronto, ON, Canada
e-mail: joshua.gans@rotman.utoronto.ca

NBER, Cambridge, MA, USA

N. Gandal

Berglas School of Economics, Tel Aviv University, Tel Aviv, Israel
e-mail: gandal@tauex.tau.ac.il

This chapter examines the underlying way that blockchain technologies achieve an ‘immutable record’ of past performance. In so doing, we focus on the dimensions and choice in the economic design of distributed ledger systems and examine their robustness using the tools of game theory. As such, we do not examine the demand for verification per se and instead focus on how blockchain technologies can supply such verification.

What do we mean when we say, ‘blockchain technologies’? While the blockchain itself is often associated with the innovations of Satoshi Nakamoto (2008), the history really begins with distributed ledger technologies. In particular, Haber and Stornetta (1990) tackled the question of how to timestamp a digital document. They noted that one way this could be done is to send the document to a time stamping service (or TSS) that would record the document (or more specifically a ‘hash’ of the document which, as will be explained below, allows it later to be verified) and then, at a later date, the TSS could verify to any party that the date the document was provided. However, this solution required all parties to trust the TSS. Cryptography would not be sufficient to secure this trust. Haber and Stornetta then argued that the trusted third party could be eliminated from the equation. Instead of the record being made with a centralized service, at the time the stamp is created, it is recorded on a ledger that is distributed among, what we will term here, ‘nodes.’ In this situation, when someone wants to verify the time stamp of a particular digital document, they can communicate with one or more of the nodes for that verification. Haber and Stornetta then went on to demonstrate the practicality of their solution by publishing a hash of their ledger each week in the *New York Times*. The hash is a unique ID that can only be recreated if you have the original records. By publishing the hash, then every, still existing *Times* edition is a distributed record of the ledger. Specifically, to change the ledger after it is recorded in the *Times*, requires all of the copies of the *Times* to be changed.

Significantly, Haber and Stornetta did not just publish a hash of the entries that they had received that week in the *Times*. Instead, each group—or as we call them now, block—of entries was hashed along with the hash of the previous block of entries. This formed a chain. In other words, in order to change an entry from 1992, you would not only have to change the record from that time but from all future times. Suffice it to say, tampering with the blockchain would be seemingly impossible. The Haber and Stornetta blockchain has been operating for almost three decades.¹

¹ See Oberhaus, Daniel, “The World’s Oldest Blockchain Has Been Hiding in the *New York Times* Since 1995,” *Vice*, available at https://www.vice.com/en_us/article/j5nzx4/what-was-the-first-blockchain.

The proposal for Bitcoin outlined by Nakamoto (2008) took this basic idea and scaled it in a way that would allow it to handle the speed and transaction volume that would be required for a network of digital payments. Bitcoin would be a ledger that recorded the ownership of digital assets (called bitcoins) the supply of which would be regulated by the protocol. At any given time, the ledger would identify the ownership (or technically the public key) associated with each bitcoin (or fraction of a bitcoin). Thus, if someone wanted to offer the transfer of ownership of a bitcoin as payment for some other service, they would need only verify their ownership and then send a message to the network to transfer that ownership to another user. These messages would then be bundled into blocks of transactions.

This led to the second new element in the Bitcoin blockchain—the consensus mechanism. How do nodes in the network ‘agree’ to append a block of transactions to the blockchain as part of the immutable record? Nakamoto outlined what is termed a ‘proof of work’ scheme. In this situation, nodes would compete in a game to solve a computational puzzle and the winner would earn both a reward and the right to propose the next block to the chain. This mechanism would impose a cost (in terms of the hardware and energy required) on earning the right to propose a block. As we will discuss below, the idea is that this cost would limit the ability of ‘bad actors’ to propose blocks that altered the record.

Since Bitcoin, the development of blockchains has moved beyond time stamping of documents and recording of transactions and now includes broader applications of verification including the provision of automated contracts and other decision-mechanisms. In addition, alternative designs for achieving consensus on digital ledgers—most notably, proof of stake—have been developed. These developments have been aimed at improving the speed, scale, resource use, and complexity that can be achieved by blockchain technologies.

Here we will focus on the consensus mechanisms, their economic operation, their costs, and their sustainability. The remainder of the chapter proceeds as follows: Sect. 2 introduces and discusses the four key elements of blockchain design. Section 3 examines equilibrium with honest actors in a setting where there is free entry into the network. Section 4 examines equilibrium with potentially dishonest actors. This section presents a condition that ensures that actors will be honest in equilibrium, that is, the condition insures incentive compatibility. In Sect. 5, we derive the efficient design that satisfies both the free entry and incentive compatible constraints. Sections 2–5 employ the ‘proof of work’ consensus mechanism. Section 6 expands the analysis to a ‘proof of stake’ consensus mechanism. In Sect. 7, we examine

whether coordination on the longest chain is a Nash equilibrium. Section 8 offers brief concluding remarks.

2 The Dimensions of Design

Nodes in blockchain networks are responsible for validating and appending new blocks to any chain. Validating transactions is a potentially costly activity as are the storage and access costs that allow users to use the blockchain to verify past information. For this reason, nodes have to be paid for their costly actions (which we denote by t). Let R be the reward for performing those actions. In order for the network to operate, it must be the case that $R \geq t$. Huberman et al. (2019) model the way in which transaction fees arise in equilibrium in the Bitcoin protocol to ensure that nodes are compensated for the costs associated with transaction processing. They show that, in equilibrium, there are some fundamental limits to the scale of the network in processing transactions. That said, for reasons that will become apparent, we will not focus on the costs of processing transactions here and ignore t in what follows.

The main challenge in the design of blockchain networks is that nodes cannot be trusted not to tamper with the blockchain. In other words, the consensus that arises is not simply something that might result because there are errors in the ways that messages are sent and received across the network but because some agents may deliberately provide false transaction records in an attempt to have them embedded as part of the record. Thus, the rewards to node operators will not only cover the costs of processing records but also have to ensure that the incentives of ‘bad’ actors are muted.

The mechanism that is used is one of ‘random selection.’ In this mechanism, each node is given a chance of becoming the node that proposes a block and receives a reward for so doing. While the reward might be anything, typically it is in the form of tokens from the network (e.g., for Bitcoin, it is a specified number of bitcoins); that is, if P is the number of tokens awarded and e is the dollar to token exchange rate, then $R = eP$. If p_i is the probability of node i receiving the reward and c_i are node i ’s costs, then the payoff to i is $p_i eP - c_i$. If this is positive, node i will operate, otherwise it will not.

There are four key design elements that determine the payoff to a node.

1. **Requirements:** who can become a node?
2. **Selection:** what determines the probability, p_i , that a given node i has in becoming the selected node?

3. **Performance:** what tasks do nodes perform and how do these impact on their costs, c_i ?
4. **Reward:** how many tokens, P , are awarded to the selected node?

We discuss each in turn.

In terms of **requirements** to become a node, different networks operate according to different policies. At one end of the spectrum are networks such as that proposed by the Libra Association whereby nodes will be appointed and vetted by the association itself as to be trusted. At the other end of the spectrum are networks such as Bitcoin and Ethereum whereby anyone, anywhere can become a node in the network. This is typically called a *permissionless* network as there is no entity that censors who can participate in the network.

Permissionless networks mean that the network agrees to not consider the identity of a node or its ownership in determining whether a node is selected to process a block. To achieve this a selection probability must have several properties as identified by Leshno and Strack (2019). One property is anonymity in that any two nodes can change their identities and inherit the selection probability of the other. Importantly, it does not allow the protocol to condition the history of the node's behavior. Another property is that a node cannot split its performance into two or more entities and pose as a new entrant to increase their selection probability—that is, they cannot engage in what are called Sybil attacks. It is this that ensures that free entry is possible and insiders cannot undertake certain actions that prevent others from entering. A final property is to ensure that insiders cannot increase their selection probability by merging. In other words, a permissionless network must forestall any incumbent advantages to effectively ensure that anyone can participate in the network on equal terms.

Leshno and Strack (2019) show that these three properties impose important constraints on the **selection** mechanism. Suppose that performance of node i is measured by a real variable, x_i (e.g., the number of computations performed). Then the only selection mechanism that respects the three properties above is a proportional rule:

$$p_i = \frac{x_i}{\sum_{j=1}^N x_j}$$

This rule makes the operation of the proof of work contest operating like an all-pay auction. A node makes a choice of x_i at a cost of $c_i(x_i)$. Holding the

contributions of other nodes (and potential nodes) as given, i makes a choice x_i^* that satisfies:

$$eP \sum_{j \neq i} p_j = c'_i(x_i^*) \sum_{j=1}^N x_j$$

Thus, the marginal cost of the node is set equal to the ratio of the probability of not being selected relative to the total performance of the network.

The significant thing about this proportional rule is that, as Ma et al. (2019) show, it is satisfied by the Bitcoin protocol. **Performance** in the Bitcoin network involves participating in a computational game. In the Bitcoin network, a node will be selected to process a block of transactions if it completes a cryptographic puzzle before any other node. The puzzle is based on Hashcash, the proof of work proposed by Back (2002) as a way of preventing spam email. The nodes search for what is called a hash function that takes an input of some arbitrary length and maps it to some output of a fixed length (i.e., the hash). This hash function is difficult and costly to compute but can be easily verified once found. The likelihood of locating it is uniformly distributed. Finally, it is practically impossible to invert so there are no shortcuts. Nodes search for a hash function based on the previous hash of the blockchain that involves a string followed by a certain number of zeros. It is the number of zeros (K) that determines how computationally difficult the problem is. Each node must search the space of possible hash functions to find the right one. This is a brute force but means that the probability that a given node finds the correct function follows a Poisson process. Indeed, the underlying game becomes very familiar in economics as it is a memoryless stochastic game that forms the basis of many models of patent races. Thus, one can easily compute the expected time a node that conducts x_i computations in a period will likely solve a puzzle of a given level of difficulty. That time is distributed Gamma(K, x_i). As Ma et al. (2019) show, this has the same equilibrium properties as emerges in the patent race literature. The main complication that arises beyond that literature is that K (the computational difficulty) is endogenous to the protocol. In Bitcoin, the computational difficulty is adjusted periodically to ensure that the total time for completion of the cryptographic puzzle by the 'system' is 10 minutes. That is, as more nodes devote resources to computations, the system adjusts the difficulty of the puzzle to slowdown the rate at which it is solved. Nonetheless, they show that the end result is that the contest to mine each block involves proportional selection. Specifically, N symmetric nodes will have the same probability ($1/N$) of being selected.

Given this, the amount of computational resources devoted to the network will be determined by the **reward** (eP) associated with being selected. While the exchange rate (e) is determined by broad market forces, the amount of tokens that is awarded to the selected node is a design decision. In Bitcoin, tokens were only issued in the course of participating in the computational game. Initially, the number of tokens awarded was 50 bitcoin but for every 210,000 blocks the reward is halved. As of 2019, it is 12.5 bitcoin. As the reward ratchets downward, the total amount of computation performed is reduced accordingly. The procedure was designed so that at most 21 million bitcoins will be mined putting a cap on the future supply of money. However, as time progresses, should bitcoins be used for transactions, transaction fees will become the more important part of rewards for nodes. At present, because rewards are primarily made up of tokens, it is the value of those tokens externally that drives the incentives of nodes to devote resources to computation.

In summary, any given proof of work network has several design decisions that it can make that determine the payoffs to node providers. First, they can change the requirements as to who becomes a node and, in some situations, might choose to limit the number of nodes operating on the network (N) at any given time. Second, they can change the selection mechanism although as already stated, for permissionless networks, that form of that mechanism is somewhat constrained. Third, they can make decisions on the cost and amount of resources required by nodes (that is, c_i). This is done by changing the nature of the computational difficulty associated with proof of work. This can have impact on the technology of computation. For instance, Bitcoin mining has evolved to use specialized ASIC chips whereas Ethereum's computational puzzle does not confer an advantage to such specialization. As will be discussed below, this can impact the economic performance and sustainability of the system. Finally, the block reward (P) can be chosen. Below we will see how these choices interact to determine the overall efficiency of blockchain networks.

3 Equilibrium with 'Honest' Nodes

Before turning to consider what happens if there are 'bad' or 'dishonest' actors present who might become nodes for nefarious purposes, it is useful to consider a 'honest' equilibrium where nodes are concerned only about their expected payoff in a simple sense. To see this, suppose that all nodes are symmetric in their size and costs (i.e., let $c_i(x_i^*) = c$; a constant) implying

that each has a probability ($1/N$) of winning the block reward (eP). The participation constraint for nodes is, therefore:

$$\frac{1}{N}eP \geq c \quad (\text{PC})$$

If this constraint were not satisfied, no one would choose to become a node.

One issue that is often discussed is that proof of work schemes require significant resource usage—particularly in terms of energy—approximately Nc . Thus, one goal of a network designer would be to reduce those costs. As Ma et al. (2019) show, in the Bitcoin protocol, if the number of nodes was regulated to be $N = 1$, then that node would act in a manner that ensures that the computational difficulty of the puzzle would be as low as possible. In this case, Nc could be arbitrarily small. This illustrates one of the benefits of having a non-decentralized network—it can operate more efficiently. Of course, this would mean that the network was not operating in a trustless manner.

By contrast, in a permissionless network, free entry will ensure that the number of nodes (\hat{N}) will be such that:

$$\frac{1}{\hat{N}}eP = c \quad (\text{FE})$$

In this case, total resource use would be $\hat{N}c = eP$. Note, however, that by changing the computational difficulty (c) there would no impact on total resource use as the number of nodes would adjust accordingly. Thus, it is only by changing the block reward, that total resource use will adjust. Of course, this might have an impact on the exchange rate (e) as well.

4 Equilibrium with (Potentially) Dishonest Nodes

The (FE) condition dictates what drives nodes to enter when they are ‘honest’ in the sense of being interested in processing transactions and validating blocks. However, nodes could also be ‘dishonest’ in the sense of having other goals that cause them to want to append blocks with information they know to be false (e.g., as might arise in a double spend attack or in an attempt to sabotage the network for other reasons). A sustainable blockchain has to be robust against such agents; deterring their entry.

In examining this, we follow Budish's (2018) investigation of proof of work protocols. Assume that there are N honest miners. Conducting activities that are dishonest requires effective control of the network. At a minimum, in say, the Bitcoin protocol, this requires a dishonest node to control a majority of computing capacity—specifically, they need to add computing power equivalent to $N + \varepsilon$. This means that the cost of conducting dishonest activities on the network is at least Nc per block.

This type of majority attack allows one node to gain control of the network for a period of time. This may be, for instance, to engage in double spending. In that situation, the attacker conducts a transaction at some point which is settled on the blockchain and in the real world. During that period of time, the attacker launches their own private network and conducts proof of work on it. As that network still adheres to the proof of work protocol, mining it costs more than mining the actual (public) network. However, during that period, the attacker earns the block reward on the private network. If the attacker has greater computational power than the public network, eventually it will have a network with the longest chain. As will be explained below, the length of the blockchain is the primary means by which nodes coordinate their mining activities. Once the private network has the longest chain, the attacker makes it public attracting other nodes to work on it. That new network potentially reverses the previous transaction and allows the attacker to spend those bitcoins again. The party to the previous transaction loses out. It is the fact of more computing power being applied to the private network by a dishonest miner that means that it will eventually generate the longest chain and be able to 'infect' the primary blockchain with its fork. This procedure is a precursor to a dishonest miner engaging in multiple spends on cryptocurrency. Of course, the double spend is potentially discoverable so such an attack may irreparably harm the network.

Budish (2018) considers two limiting factors on a simple majority attack. First, some activities from dishonest miners may require more than a simple majority to implement. For instance, control to achieve a fork may require control of $\frac{A}{A+1}$ percent of the nodes. Thus, the cost per block for entry by a dishonest miner would be ANc (where $A > 1$).

Second, for some activities that involve interaction outside the blockchain (such as a multi-spend attack), control of the blockchain cannot be confined to just the block in question but may require a time period to elapse. Thus, the dishonest node may have to control the network for a time which translates into adding t blocks.

Offsetting these limiting factors is the fact noted earlier that, while controlling a network, a block reward (eP) will be earned for each block added. That

reward accrues to the dishonest miner. Putting these together, the net cost to the dishonest node is $(ANc - eP)t$. The entry decision of a dishonest node will be driven by the benefits they receive from such control—that is, from dishonest activities. Suppose that the private benefit for an attacker is $V(e)$.

Given these costs and benefits to a dishonest node, we can see that dishonest entry to a blockchain network consisting of N nodes will not be profitable if:

$$AtNc - teP \geq V(e) \Rightarrow Nc \geq \frac{V(e) + teP}{At} \quad (\text{IC})$$

This incentive compatibility condition comes from Budish (2018). The left-hand side $(AtNc - teP)$ is the cost of controlling NA nodes for t periods less the block reward earned during the control period. The right-hand side, $V(e)$, is the benefit of exercising that control for personal benefit. In the analysis that follows, we suppose that the private benefit for an attacker is $V(e)$; a non-decreasing function; that is, the more valuable is cryptocurrency, the greater is the private benefit from dishonest activities.

For an equilibrium to exist that allows for the blockchain to operate both the (PC) and (IC) conditions must be met. The (PC) says that the total network costs (Nc) must not be too high while the (IC) says that they should not be too low. Putting these together we have:

$$eP \geq Nc \geq \frac{V(e) + teP}{At}$$

Thus, a necessary condition for a blockchain to be sustained in equilibrium where there are potentially dishonest nodes is that:

$$eP \geq \frac{V(e) + teP}{At} \Rightarrow eP(A - 1)t \geq V(e) \quad (\text{B})$$

This condition was derived by Budish (2018)—and we call it the Budish condition. If this condition holds then, even if the (PC) condition binds (as it would in a permissionless network), then the (IC) condition will be satisfied and not bind.

The attack contemplated here is less costly for the attacker because it is what Budish (2018) calls a ‘flow’ rather than a ‘stock’ attack. Specifically, while the private chain is in operation, the attacker earns a block reward but more importantly they are only expending energy in doing so. If that energy cost is c_E per node and the rental cost of computer equipment is rC per node, then $c = rC + c_E$. The idea is that once the attack stops, the attacker can stop

using energy and can repurpose (or resell) the computer hardware. This latter option was, in fact, a feature of the Bitcoin system as proposed by Nakamoto (2008) in that it would democratize participation in the network by lowering the cost of entering to operate a node; the so-called ‘one CPU-one-vote’ ideal.

As Bitcoin has evolved, nodes have adopted technologies to give them an advantage in solving the cryptographic puzzle. One of these was the use of ASICs (application-specific integrated circuits). These could be manufactured for the specific purpose of Bitcoin mining but this also meant that they would not be repurposable. In this situation, an attacker would not be able to recoup the hardware cost following the attack and that cost would be sunk. This would make the attack a ‘stock’ attack and the cost of the attack would be $N(C + c_E)$. The (IC) condition would then become:

$$AN(1 - r)C + AtN(rC + c_E) - teP \geq V(e)$$

$$\Rightarrow Nc \geq \frac{V(e) - AN(1 - r)C + teP}{At}$$

This would allow for sustainability against a higher $V(e)$ than the ‘flow’ (IC) constraint. Nonetheless, the analysis here would proceed substituting $V(e) - AN(1 - r)C$ or $V(e)$.

That said, there are three reasons why this ‘stock’ analysis may not be the appropriate one. First, for some blockchains (including Ethereum until recently), ASIC chips do not confer an advantage. Second, in some situations, it may be more efficient for the attacker to expand the number of nodes in the attack and use repurposable chips rather than ASIC chips in which case the ‘flow’ analysis applies. Finally, it may be that an attack would not make the ASIC chips unusable because the network is not significantly damaged by the attack. However, this last option likely depends on the purpose of the attack, the possibility of future attacks, and other real-world specifics of the blockchain in question.

5 Efficient Design

Having identified the participation and incentive compatibility constraints that will allow for a sustainable blockchain, we can now consider the choice of design parameters that will minimize resource use subject to those constraints. That is, what are the choices of recruitment (N), performance (c), and reward (P) that solve:

$$\text{Min}_{c,N,P} cN \text{ subject to } ANtc - teP \geq V(e) \text{ and } Nc \leq eP$$

Gans and Gandal (2019) note that, when it binds, the (IC) constraint implies that $Nc = \frac{V(e) + teP}{At}$. Observe that by decreasing the reward (P), it is possible to maintain incentives at the expense of fewer resources. What prevents this from being lowered all the way to zero is, of course, the (PC) condition that requires that there is a sufficient reward that the nodes are not making losses. The minimum possible P is, therefore, $\frac{1}{e}Nc$ which results in $Nc = eP$. Given this, the P that minimizes resource cost satisfies:

$$eP = \frac{V(e) + teP}{At} \Rightarrow P = \frac{V(e)}{e(A-1)t}$$

By setting P at this level, the (PC) and (IC) constraints both bind and the cost (Nc) is minimized at $Nc = \frac{V(e)}{(A-1)t}$.

There are several things to note about this outcome. First, if P is chosen optimally, then the choices of N and c are determined. In other words, neither the difficulty of performance (c) nor whether the number of nodes is fixed or left to free entry (as in a permissionless network) matters. In other words, for a sustainable blockchain, resource use is not directly impacted on by the number of nodes nor their proof of work cost. Instead, it is the reward function that matters.

Second, the reward function optimally varies with the exchange rate (e). This is not something that any blockchain protocol (including Bitcoin) does at present. However, to create a sustainable blockchain, adjustment to ‘market conditions’ is a useful property for the block reward.

Therefore, the block reward should change as the exchange rate, e , changes. But in which direction should the change go? Note that, as e increases, the impact on sustainability depends on the relationship between e and V . If V is non-decreasing in e , the effect of an increase in e depends upon the sign of $(eV'(e) - V(e))/e^2$ or whether the elasticity of V with respect to e is greater than 1. In this case, a higher value for the cryptocurrency increases the probability that the blockchain will be vulnerable to attack. Therefore, if e rises, to ensure that the blockchain is sustainable, P has to rise as well if the elasticity of V with respect to e (that is, $\frac{e}{V(e)}V'(e)$) is greater than unity implying that an increase in e causes a higher proportionate increase in V .

To put this another way, for Bitcoin, as noted earlier, P is set to fall over time as part of the protocol while e is expected to be higher (at least by Bitcoin enthusiasts). The only way these two changes would imply long-term sustainability of the blockchain is if $\frac{e}{V(e)}V'(e) < 1$ or that the private payoff from an attack rises proportionately less than an increase in the value of bitcoin. Note that, for multiple-spend attacks, this may be a reasonable assumption

as the scale of such attacks is capped by real-world motivations.² For other issues (such as sabotage of the network), that may not be the case and Bitcoin has to be robust to them all. Without more knowledge regarding V , it is not possible to assess whether these changes in the block reward make Bitcoin more or less sustainable.

6 Other Consensus Mechanisms

The largest blockchain networks (Bitcoin and Ethereum) rely on a Proof of Work (PoW) scheme such as those analyzed above to achieve consensus. The energy cost associated with these has drawn criticism so there is increased interest in alternatives. For instance, Ethereum are looking toward a Proof of Stake (PoS) protocol where nodes hold an amount of tokens to align their incentives and a related notion is Proof of Burn where nodes publicly ‘burn’ tokens (by sending them to invalid public addresses) although that may impact upon the monetary policy of a network (see Saleh 2018).³ Here we outline PoS and, following Gans and Gandal (2019), show how it impacts the above analysis.⁴

PoS achieves consensus on blocks by requiring nodes to *stake* a sufficient quantity of tokens in order to be considered as a validator for a new block of transactions. There are, however, different ways in which validator nodes are selected but in chain-based methods, the selection process is similar to PoW where each node has a proportionate probability of being selected to propose a block and we will focus on that here.

Suppose that S tokens are required for a stake and that the dollar interest rate is r , then (in terms of our PoW notation) $c = reS$; that is, the opportunity cost per period of resources a node must hold to be ‘staked’ (e.g., the lost earnings on fiat currency that is held in tokens). The stake, S , can be chosen

² For instance, for a double-spend attack, the value of a transaction will drive incentives to attack the network and, in effect, defraud a ‘real-world’ payor. For instance, at a Bitcoin exchange rate of \$10,000 ($= e$), a simple majority value for A ($= 1.01$), and a block reward of 12.5 BTC, then, if the required time is 6 hours (so assuming 10 minutes per block, $t = 36$), the value (V) must be greater than \$45,000 to be worthwhile. If the required time is just 1 hour ($t = 6$) that threshold falls to \$7,500. It is this type of calculation that caused Budish (2018) to conclude that Bitcoin may not be sustainable.

³ In closed and small networks, computer scientists have also applied Byzantine Fault Tolerance as a means of achieving consensus. For a review see Amoussou-Guenou, Biais, Potop-Butucaru, and Tucci-Piergiovanni (2019).

⁴ See also Saleh (2019) for an examination and debunking of the so-called ‘nothing at stake’ critique of PoS.

in a permissionless PoS protocol which means that, potentially like c , it is a design decision.

This allows us to write a (FE) condition for a PoS network. Like in a PoW network, any node has a probability $1/N$ of earning a block reward. Thus, the expected per block payoff to a node is $\frac{1}{N}eP - reS$. Thus, the (FE) condition becomes (ignoring integer constraints):

$$S = \frac{P}{rN} \quad (\text{FE})'$$

Note that, unlike PoW, the (FE) condition does not depend on the exchange value (e) of tokens. However, like PoW, the total network ‘cost’ (that is, SrN) is fixed—that is, an increase in S , causes a proportionate reduction in the number of nodes (N).

What about protection against attacks by dishonest nodes? Both PoS methods are vulnerable to attack forms based on establishing a private chain with altered transactions before releasing to publicly. With PoW, this entails a cost as a dishonest node is required to perform the PoW of the entire network in order to obtain the longest chain upon publication. With PoS, there is no such cost. The main challenge comes, however, that when the alternative chain is published there is the challenge of getting other nodes to accept it. For nodes that were online while the alternative chain was being written, they will be able to identify the alternative chain. For new nodes or ones that were offline, they cannot tell which is the legitimate chain. Thus, for an attack to be successful, the dishonest node needs to take actions that would shift the share of online versus other nodes. We assume that this takes time (t periods.)

Such attacks rely on the attacker building on both the main chain and their alternative at the same time. This is something that is possible with PoS but costly for PoW. However, PoS networks have implemented various methods to guard against this. One such method is called ‘slashing.’ This involves the stake of a node being reduced or destroyed if it is found that they have worked on multiple chains. This is something that can be algorithmically detected. Importantly, when slashing is effective, it turns a ‘flow’ attack into a ‘stock’ attack as the stake itself can be taken from the attacker should they be discovered.

That said, while slashing can prevent ‘low scale’ attacks on the network, PoS networks are still vulnerable to a majority attack—as we examined for PoW. Such an attack requires the attacker to stake a supermajority of nodes for t periods. If the value of an attack is $V(e)$ as before, then the cost of the attack is $ANtreS$ less than the block reward teP earned on the alternative chain. Note that block reward accrues to the attacker precisely because

slashing or other mechanisms penalizes others if they work on the alternative chain leaving all of the block rewards to the attacker.

Given this, the (IC) constraint is

$$ANteSr - teP \geq V(e) \Rightarrow S \geq \frac{V(e) + teP}{eANtr} \quad (\text{IC})'$$

Thus, so long as the stake, S , is sufficiently high, an attack can be prevented.

We can now perform a comparable exercise to that for PoW to examine what will determine the sustainability of a PoS blockchain. In particular, using (FE)' and (IC)', if the (IC)' constraint is to be satisfied while the equilibrium number of nodes is determined by the (FE)' condition that requires:

$$(A - 1)teP \geq V(e) \quad (\text{B})'$$

Note, critically, that this condition is *identical* to the PoW Budish condition (B). Moreover, it is independent of the level of the stake (S).

In other words, despite the ability to control S , there are no design choices under proof of stake that will lead to greater sustainability than under proof of work. In fact, in designing a permissionless blockchain (even though S can be chosen), the minimum block reward that will create a sustainable outcome is $P = \frac{V(e)}{e(A-1)t}$. This also means that the same elasticity condition on $V(e)$ drives whether, for a fixed block reward, the network will be more sustainable as e grows.

It is useful to note, however, that the mechanism for sustainability is different. In particular, under PoS, the FE condition is independent of the exchange rate (e). In other words, the number of nodes will not change as the exchange rate changes and will be pinned down by the level of the stake. Thus, the size of the permissionless network can be controlled by changing the stake. This is not the case in PoW, since given P and e , c (which is a design variable) determines the network size (N) in the permissionless PoW mechanism.

7 Coordination

Thus far, we have analyzed the basic underlying choices of nodes in appending blocks to the blockchain. However, the final aspect of achieving consensus is with regard to the overall chain that becomes the locus of activity.

In particular, as blockchains are open source, it is possible for anyone to ‘fork’ the chain with their own rules—both technical and otherwise—and, if they can achieve enough interest, end up competing with or displacing the incumbent chain (Barrera and Hurder 2018). Importantly, in the Bitcoin protocol there is no rule regarding which chain nodes may work on. Nakamoto realized this and proposed that when appending blocks, the ‘longest chain rule’ should be applied—that is, nodes should work on whatever blockchain has the most blocks. It was this rule that was potentially exploited in the majority attack scenarios considered earlier. But this rule has the appealing property that the longest chain in a PoW system is likely to be the one where the incumbent nodes have the greatest interest. Thus, in order to enact changes, a party would have to convince a (computational) majority of nodes; something that would be hard if a chain is purely fraudulent or possessed code that would allow someone to take control of the network.

Given this, there have been numerous investigations into coordination and whether it is a Nash equilibrium for the longest chain ‘norm’ to persist (Kroll et al. 2013). The most complete treatment is Biais et al. (2019). They prove a folk theorem that demonstrates that the ‘longest chain rule’ is a dynamic Nash equilibrium. However, there are other possibilities. They show that nodes do not need a (computational) majority to successfully attack the system and render it unstable. In their model, forks can delay the achievement of consensus and reduce the flexibility in upgrading the system and resolving problems. Such aspects have been noted by Arruñada and Garicano (2018) who argue that ‘within protocol’ governance is unlikely and, at the very least, soft power and relationships will guide the evolution of blockchain networks.

In the end, there are costs of achieving consensus and coordinating on a chain, especially one whose ledger is ‘correct.’ Abadi and Brunnermeier (2018) model these costs and show that the costs of operating a blockchain that is both decentralized and correct are necessarily higher in order for nodes to be appropriately incentivized. Attempts to reduce those costs necessarily make the blockchain vulnerable to misreporting while decentralization is not compatible with reducing costs and ensuring correct reporting. In other words, there is a trilemma in achieving a blockchain that is correct, decentralized, and low cost. Only two out of three of these goals can be achieved. Suffice it to say, these costs have already emerged in decentralized blockchains and it remains an open question whether a blockchain that is more efficient and operates at the scale of networks such as Visa and Mastercard is possible.

8 Brief Concluding Remarks

In this chapter, we examined the equilibrium properties of consensus mechanisms for the blockchain under both proof of work and proof of stake when the blockchain is permissionless, i.e., when there is free entry into becoming a node. It is straightforward to extend the analysis to permissioned blockchains, in which the blockchain itself can determine who can become a node.⁵

Acknowledgements We are grateful to Orbs for an academic grant and to Eric Budish and Oded Noam for helpful discussion, suggestions, and comments that significantly improved the paper. All opinions (to the extent that they are opinions) are those of the authors.

References

- Abadi, Joseph, and Markus Brunnermeier. 2018. Blockchain economics, *mimeo.*, Princeton.
- Amoussou-Guenou, Yackolley, Bruno Biais, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. 2019. Rationals vs byzantines in consensus-based blockchains, arXiv preprint [arXiv:1902.07895](https://arxiv.org/abs/1902.07895).
- Arruñada, B. and Luis Garicano. 2018. Blockchain: The birth of decentralized governance, *mimeo.*, IESE.
- Back, A. 2002. Hashcash—A denial of service counter-measure. Available at: <https://www.hashcash.org/papers/hashcash.pdf>.
- Barrera, C., and S. Hurder. 2018, June 16. Blockchain upgrade as a coordination game. Available at SSRN: <https://ssrn.com/abstract=3192208>.
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta. 2019. The blockchain folk theorem. *the Review of Financial Studies* 32: 1662–1715.
- Budish, E. 2018. The economic limits of bitcoin and the blockchain. Working paper. Available at: <https://faculty.chicagobooth.edu/eric.budish/research/Economic-Limits-Bitcoin-Blockchain.pdf>.
- Catalini, C., and J.S. Gans. 2019, Forthcoming. Some simple economics of the blockchain. *Proceedings of the ACM*.
- Gans, J.S., and N. Gandal. 2019. More (or less) economic limits of the blockchain. Working Paper, NBER.
- Haber, S., and W.S. Stornetta. 1990, August. How to time-stamp a digital document. In *Conference on the Theory and Application of Cryptography*, 437–455. Berlin, Heidelberg: Springer.

⁵ See Gans and Gandal (2019) for details.

- Hart, O. 1995. *Firms, contracts, and financial structure*. Clarendon Press.
- Huberman, Gur, and Jacob Leshno, and Ciamac C. Moallemi. 2019, March 13. An economic analysis of the bitcoin payment system. Columbia Business School Research Paper No. 17-92.
- Kocherlakota, N.R. 1998. Money is memory. *Journal of Economic Theory* 81 (2): 232–251.
- Kroll, Joshua A., Ian C. Davey, and Edward W. Felten. 2013. The economics of bitcoin mining, or bitcoin in the presence of adversaries. In *Proceedings of WEIS*, vol. 2013. Citeseer.
- Leshno, J., and P. Strack. 2019. Bitcoin: An impossibility theorem for proof of work protocols, *mimeo.*, Yale.
- Ma, J., J.S. Gans, and R. Tourky. 2019. Market structure in bitcoin mining, *mimeo.*, ANU.
- Nakamoto, S. 2008. *Bitcoin: A peer-to-peer electronic cash system*. Manubot.
- Saleh, F. 2018. Volatility and welfare in a crypto economy. Available at SSRN 3235467.
- Saleh, F. 2019. Blockchain without waste: Proof-of-stake. Available at SSRN 3183935.



Multi-Stakeholder Governance in the Era of Decentralized Finance

Shin'ichiro Matsuo and Ryosuke Ushida

1 Introduction

Blockchain technology, which was originally introduced by the anonymous Satoshi Nakamoto's Bitcoin paper, proposed many innovations to the world of payment, accounting, and bookkeeping. One of the main innovations that blockchain technology creates is a new form of trust model. In using blockchain technology, we can eliminate any trusted party in updating the shared ledger. Logics, algorithms, and mathematics to update the shared ledger are given by grassroots programming code. In fact, the software code of Bitcoin, which realizes a simple business logic, payment, was initiated by an anonymous person(s) and maintained by a group of grassroots engineers

S. Matsuo (✉)

Department of Computer Science, Georgetown University, Washington, DC, USA
e-mail: sm3377@georgetown.edu

NTT Research Inc., Tokyo, Japan

R. Ushida

Cyber SMART, Georgetown University, Washington, DC, USA
e-mail: ru64@georgetown.edu

Fintech and Innovation Office, Japan Financial Services Agency, Tokyo, Japan

known as the “Bitcoin Core.” The paper and software code were not peer-reviewed, but is widely deployed and is now operated by over 10,000 nodes, and has created a global payment space without a trusted party.

This phenomenon is completely different from the design of the ordinary financial industry. The software code is deployed and executed without permission, which is common among the Internet community. In general, financial regulators allow only approved businesses. However, Bitcoin created a global space outside of this. From the innovators’ perspective, permissionless innovation is essential to creating new business. However, deployment without a peer-review style is creating problems. For example, many cryptoasset exchange/custodians have been attacked and huge amounts of cryptoassets have been stolen due to careless implementation and operation. This type of incident is a serious problem for consumer/investor protection. Another problem for regulators is the lack of transparency. Many blockchain startups publish a white paper that describes their business model and overview of technology. However, in many cases, such a white paper does not describe a detailed system design suitable for a review by a third party. It is not easy for average investors and regulators to conduct due diligence of the programming code in the GitHub repository. There have many scams in the Initial Coin Offering (ICO) projects. Such problems are caused by a lack of transparency.

From a regulatory point of view, the ordinary style of oversight does not work for permissionless innovation. Hence, we need to consider a new style of governance to achieve both facilitating new innovations and achieving regulatory goals.

2 The Implications of Permissionless Blockchain for Regulators

2.1 Ordinary Regulatory Goals

Generally, the reason why regulators exist and work is to prevent market failure, which are asymmetric information, negative externality, public goods, imperfect competition, and irrational behaviors. Financial regulators exist to address these market failures. Each financial regulator articulates its mission and goals in varying parlance, and no formally agreed list of global regulatory goals exists. However, Armour et al. summarize financial regulatory goals as (1) investor and consumer protection, (2) financial stability, (3) market integrity, and (4) financial crime prevention (John Armour 2016).

2.1.1 Investor and Consumer Protection

Investor and consumer protection is a primary goal for financial regulators. For investors to assess risks and make an appropriate judgment on a given investment, they need accurate information. However, due to information asymmetry, investors may face difficulties in doing so and move on. In such a situation, regulators can intervene by requiring relevant entities to accurately disclose necessary information to investors to facilitate their decision-making processes. As described in the previous section, cryptoasset exchanges/custodians have been attacked and huge amounts of cryptoassets have been stolen. For many ICO projects, average investors cannot evaluate the risks of the project. These are examples of new types of threats to investors and consumers.

2.1.2 Financial Stability

As demonstrated by the financial crisis of 2008, the failure of a financial institution can cause a greater collapse of the financial system via a domino effect. This interconnectedness among financial institutions and markets is a negative externality that financial regulators must address. The traditional regulatory approach is to impose regulations on individual financial institutions, including capital requirements, to minimize their probability of failure. In addition, deposit insurance schemes are introduced in many jurisdictions to maintain market confidence and protect depositors to prevent bank runs.

However, the financial crisis revealed that simply decreasing an individual financial institution's probability of default. Deposit insurance does not adequately protect institutional investors, such as other banks, which caused the run-like behavior in the market. In addition, it turned out that the domino effect could also be enhanced not only by the balance sheet problems but also through the deterioration in liquidity level, mispricing of financial instruments, and highly correlated investment strategies with deleverage. Thus, after the crisis, financial regulators introduced enhanced regulation on systemically important financial institutions (SIFIs) and new types of regulations called "macroprudential regulations," such as countercyclical buffer and loan-to-value ratio (LTV), which focus on the overall stability of the system.

2.1.3 Market Integrity

Both primary and secondary markets play a crucial role in capital formation, capital allocation, and liquidity provision. Without the secondary market, market prices do not reflect a real-time estimation of the valuation of a given firm. The speed and accuracy at which the market price reflects new information is called “information efficiency,” which is important to ensure market integrity. Information efficiency acts as a real-time monitoring tool for a firm’s investment policies, guiding managers of the firm. The more accurate information is made available to the secondary market, the more accurate decisions investors can make.

In this regard, market information can be considered a public good. However, asymmetric information can be a source of inefficiency, as managers tend to be reluctant to disclose information for various reasons. In addition, unfair trading, such as insider trading, could harm market integrity by impairing the quality of market information. Thus, regulators need to mandate timely, accurate, and adequate disclosure of necessary information as well as crack down on unfair trading activities.

2.1.4 Financial Crime Prevention

Armour et al. suggest that preventing financial crimes that take advantage of the financial system for socially harmful purposes (“public bad”), for example, money laundering, should be distinguished as an independent regulatory goal (John Armour 2016).

Financial regulators and other national authorities, including financial intelligence units (“FIU”), take several measures to achieve this goal. In most jurisdictions, regulators criminalize money laundering and mandate financial intermediaries such as banks as gatekeepers for the financial system to effectively enforce regulations related to criminal activities. Banks are required to conduct “know your customer” (“KYC”) screening when their customers open accounts as well as ongoing monitoring to report suspicious activities in their accounts to authorities, including FIUs. With the information provided, national agencies investigate and prosecute criminal activities. For example, according to the Financial Crimes Enforcement Network (“FinCEN”), banks in the US reported 1,116,400 suspicious activities in 2019 (FINCEN 2020).

2.2 Main Difference from the Traditional Financial Industry

While regulators need to achieve these regulatory goals regardless of the technologies used in the financial ecosystem, Yuta demonstrates that inherent characteristics of blockchain technologies make it difficult for regulators to achieve the goals through traditional regulatory approaches. We summarized the arguments made by Yuta in the following sections (Yuta Takanashi et al. 2019).

2.2.1 Preclude Intermediaries from Transactions Between Cryptoassets.

As previously discussed, users are able to transact on a peer-to-peer basis, which may not need to adhere to AML/CFT/KYC regulations and could enable criminal conduct. In relation to this point, the Financial Action Task Force (FATF), an international standard-setter for financial criminal activities, has pointed out that “[l]aw enforcement cannot target one central location or entity (administrator) for investigative or asset seizure purposes,” and “[d]ecentralized convertible virtual currencies allowing anonymous person-to-person transactions may seem to exist in a digital universe entirely outside the reach of any particular country (Financial Action Task Force 2015).”

As an example, FATF’s Recommendations 15 and 16, updated in October 2020, require originator virtual asset service providers (VASPs) to submit originator and beneficiary information to beneficiary VASPs when performing VA transfer of USD/EUR 1,000 or more (Financial Action Task Force 2012). This is an entity-based regulation similar to those applied to traditional wire transfers conducted by existing financial institutions such as banks. However, users may be able to circumvent the requirement by transferring cryptoassets from originating VASP to a non-custodial wallet, which is not subject to the regulation, and subsequently from the non-custodial wallet to beneficiary VASP. These technological characteristics generate significant challenges for regulators to ensure the enforceability of the regulations relying on intermediaries.

2.2.2 Sophisticated Anonymization Technologies Confuse the Tracing of Transactions.

To prevent individual users from conducting unlawful financial transactions, regulators need to monitor transactions, including peer-to-peer transactions, and take necessary enforcement actions. However, ensuring regulability against individual users would have limitations with respect to policing capacity. Moreover, users could take advantage of privacy-enhancing technologies that anonymize users or hide the linkage between transactions within a blockchain-based system. Various technologies such as signature aggregation, zk-SNARKs (Zcash 2019), and mixing are in development as engineering communities are very keen to enhance the degree of privacy of cryptoasset transactions. In addition, layer-2 technologies such as the Lightning Network for Bitcoin also have an aspect of strengthening anonymity as it applies onion routing to hide linkages between senders and receivers of a transaction from the public (Lightning 2019; Gary Basin 2018). Such technologies have already been deployed in various projects, notably in “privacy coins” including ZCash and Monero. Despite the projects’ limited market share, their network participants already have the ability to hide their information, and even Bitcoin is subject to greater anonymity via the aforementioned developments.

In this regard, regulators face challenges in that these technologies would further complicate enforcement actions toward individual users and render them more costly. FATF highlights that they may allow greater anonymity than traditional non-cash payment methods. Virtual currency systems can be traded on the Internet, are generally characterized by non-face-to-face customer relationships, and may permit anonymous funding (cash funding or third-party funding through virtual exchangers that do not properly identify the funding source). Although many studies on techniques to de-anonymize and/or de-link transactions as well as several AML/CFT solutions for blockchain, are ongoing, we cannot attest to their efficacy at this point in time.

2.2.3 Difficulties in Stopping a Decentralized Network and Modifying Transactional Records

In conventional financial systems, regulators have the capability to order suspension of unlawful and/or problematic services or cancel or reverse the transaction in cooperation with financial intermediaries. However, in the absence of a reliable regulated entity in a blockchain-based financial

ecosystem, it is difficult to address such situations even if regulators succeed in identifying unlawful activity. While some authorities place importance on conducting transparent, rule-based supervision focusing on ex post checks rather than ex ante guidance to provide market participants with a broad degree of freedom, such an approach may not work well in decentralized financial systems.

2.3 Multidisciplinary Background of Blockchain Technology

Blockchain technology is composed of multiple backgrounds. The immutability of the data is realized by the digital signature scheme. Assurance of the order of transactions is assured by a link of cryptographic hash value among blocks. These are based on the security of cryptographic algorithms and protocols. However, the security and reliability of blockchain technology do not rely solely on cryptography. Eliminating any trusted party and single point of failure, that is, the most essential merit of blockchain against cryptographic time-stamping, is realized by introducing a peer-to-peer (P2P) network, a consensus algorithm such as proof of work (PoW) and proof of stake (PoS), and incentive mechanisms as rewarding cryptoassets. The design of a P2P network is based on network theory, and the design of a consensus algorithm is based on the theory of distributed computing. The design of incentive mechanisms is based on game theory and economics. When we consider only such technology aspects, the necessary background to obtain reliable technology is very wide. Moreover, when we apply the permissionless blockchain technology to a real-life financial business, we need to assess if the technology does not infringe on regulatory goals. Deep knowledge about laws and regulations and the capacity to communicate with regulators are also required. From the above, when we design a financial business and information system based on permissionless blockchain technology, we need a very wide range of expertise and backgrounds that are not covered by a single person and stakeholder.

2.4 Who Are Creators of the Order of Finance?

Before Bitcoin, the financial system and financial products were designed by experts in the financial industry, which has sufficient communication with financial regulators to achieve regulatory goals. Programming codes that

realize such a design of financial systems and financial products result from such a top-down style design.

However, the invention of Bitcoin introduced an important concept of “permissionless innovation” to the financial industry. This concept allows anyone to be an entity who can create a financial application without permission. This change is similar to the impact of the Internet on the communication industry. Before the Internet, a person who wanted to circulate some information to the public had to request some authority to circulate it. Examples of such authorities are newspapers, broadcast companies, and telecommunication companies. In France, Minitel was an example of a multimedia service to circulate information, and the CAPTAIN system played the same role in Japan. In such systems, permission is needed to circulate the information, that is, creating new services from grassroots was very difficult.

The Internet and the World Wide Web changed the creators in the communication ecosystem. Anyone can initiate a web server and website, and then create a new service. Many innovative Internet services have been created by people in cooperation with engineers and computer experts. This kind of “permissionless innovation” is a good solution to the innovation dilemma.

Lawrence Lessig proposed a framework of regulations in the era of open networks and programming codes (Lawrence Lessig 1999). The framework represents the relationship among factors to form a new order. That is, in the era of permissionless innovation, the order is formed by law, market, social norms, and architecture. Programming code is a part of the architecture in this context. This means that law is not a perfect tool to decide the regulatory framework in the era of permissionless innovation based on programming code.

The same situation has occurred in the area of finance since the invention of Bitcoin. The mathematics and programming code of Bitcoin implies a new order of decentralized finance. A smart contract based on Ethereum or other blockchain technology is a new platform for permissionless financial innovation. In the ordinary financial industry, the order is formed by a regulator and financial business entity, but after Bitcoin, anyone with engineering capacity can be a part of the stakeholder to form the order.

3 Stakeholders and Pain Points

3.1 Stakeholders in the Blockchain Ecosystem and Pain Points

There are four types of stakeholders in this sphere: (1) Regulator, (2) open source style engineers, (3) blockchain business entities, and (4) consumers. The regulator does not share a common language to communicate with open source style engineers, who sometimes do not want to talk with regulators; thus, there is no communication channel between them. It is hard to regulate such engineers because they are members of a community, and not a legal entity. Blockchain business entities usually want to avoid friction with regulators and start a business before the technology is matured. Consumers are not certain about the reliability of services provided by these entities due to lack of transparency. Overall, there is no one place where stakeholders can communicate based on the same language.

4 Lessons Learned from Internet Development

4.1 Global vs International

One special aspect of a permissionless blockchain such as Bitcoin is the creation of a global financial ecosystem. The term “international” represents the relationship among nations, but the term “global” represents something monolithic, thus, it is independent from nations. This difference is essential when considering the governance of the permissionless blockchain. The Internet is a global network, and no nations control its operation; it is independent from any nation. Existing financial regulations are decided by a group of nations, such as G20. However, for Bitcoin and permissionless blockchain, this framework sometimes does not work due to its global nature.

4.2 Overview of Internet Governance

Looking back at the history of cyberspace, we see that the Internet could be a good reference point to find clues to address various issues in blockchain-based financial ecosystems described in the previous chapter. This is because many challenges stemming from decentralized architecture have been reconciled through collective efforts made by a wide range of stakeholders in the

Internet community. As blockchain-based systems have a lot in common with the Internet architecture, we should learn a lesson from such efforts to develop new types of orders for the financial system of the future.

In this section, we provide a brief overview of Internet governance because it plays many important roles in addressing social problems. While the Internet is deemed one of the most successful examples of the power of decentralized community-building, it is argued that the Internet community has taken a somewhat roundabout route to becoming a global and borderless communication platform. After cumulative failures of top-down controls by jurisdictional authorities, the Internet has reached an idealistic governance model by steadily overcoming a number of societal and regulatory challenges through the gradual development and adoption of the “multi-stakeholder” approach.

The Internet has been developed primarily as an open source project through the efforts of a large collection of individuals. The Internet Society (ISOC) states that, “While the Internet evolved from a number of government-funded research projects, individuals from universities and private sector organizations led most of its early development. Since these early beginnings, management of the Internet and global Internet resources (e.g., the Domain Name System, DNS) has relied heavily on bottom-up coordination and direct participation by those interested in and impacted by related decisions.”

However, as the Internet grew and became an important social infrastructure, some government actors attempted to seize control over the Internet in a top-down manner through the United Nations. This attempt eventually failed partly because strong opposition from the Internet community and governments finally accepted the concept of “multistakeholderism” and roles of the private sector stakeholders in managing the Internet resources and a non-binding forum called the Internet Governance Forum (IGF) was established to discuss policy issues.

Currently, the Internet is managed by several different public and private multi-stakeholder entities. Some of the organizations such as IETF are working on technical development and standardization, while others such as ICANN mainly manage day-to-day operations related to critical Internet resources such as the DNS. Each of them plays different roles with distinct structures and governance mechanisms; however, they work together to guarantee that the Internet continues to function effectively as an open and consistent network facilitating global communications.

4.3 Leveraging Internet Experiences to Tackle Issues in Blockchain-Based Financial Ecosystems

As previously discussed, financial regulators face unprecedented challenges with the advent of decentralized financial technology. This implies that the traditional regulatory approach will no longer work well to achieve regulatory goals, and regulators alone may not be enough to address emerging issues. In this section, we examine how telecommunication regulators and participants in Internet governance tackle the actual problems in cyberspace and draw lessons that will provide valuable insights into the question of how to create a sound governance mechanism in the blockchain-based financial system.

The first lesson from the Internet is the importance of creating a cooperative environment among stakeholders. For example, in the case of P2P file sharing, the government failed to cooperate with the engineering community and gave false incentives to encourage development that made it even more difficult to address issues. When P2P file sharing services such as Napster became a threat to copyright holders, the governments and copyright holders mandated internet service providers (ISP) to police unlawful P2P transactions and deployed a technology called deep packet inspection (DPI) to actively monitor the flow of data and block suspicious access to their users. However, the number of infringements was not significantly reduced, partly because new decentralized technology for P2P file sharing services such as BitTorrent made it difficult for regulators to impose effective oversight. In addition, ISPs' active monitoring of their users' network communications with DPI technology has profound implications for the freedom of expression. Against this approach, there are transnational resistance movements such as A2K (access to knowledge) embracing fairness and access to knowledge. This is an example of failing to create a cooperative environment and resulting in undesirable development of the situation that we should avoid in a decentralized financial system. Unfortunately, there currently seems to be a little cooperative relationship among stakeholders in the blockchain-based financial ecosystem, which is a great risk in maintaining the healthy development of the technology and the ecosystem.

The second lesson is that regulators should make the most of the solid governance mechanisms that play a leading role in tackling issues. It is likely that governance organizations that are closer to the issue may be the best players to tackle that issue. In the case of trademark infringement in cyberspace, ICANN greatly contributed to addressing the infringement cases by establishing a globally applicable private arbitration process that enables trademark holders to challenge and recover domain name registration along

with the Uniform Domain Name Dispute Resolution Policy (UDRP). In analogy with the case of the Internet, financial regulators may wish to have private entity-driven governance mechanisms to achieve regulatory goals. However, at this stage, there is no well-established governance mechanism for blockchain-based financial ecosystems. Thus, efforts need to be made to establish a well-organized governance mechanism and develop mutual understandings among stakeholders to realize a healthy ecosystem.

The third lesson is that simply bringing together different stakeholders in one place may not be sufficient to ensure material outcomes from the activities. Although the IGF was intentionally organized as an open discussion forum for multi-stakeholder dialogues, there is no incentive for participants to follow the outcome of the discussion. In contrast, in the case of the ICANN, the Internet community follows the outcome because it manages the single point of control of the network and has a strong governance mechanism within it. Thus, one of the key implications for the governance of the blockchain-based financial ecosystem is that we need to pay particular attention to the way to provide adequate incentives or reasons for each stakeholder to abide by the outcome of the governing activities.

Finally, the legitimacy of the privately owned governance mechanism matters especially when facing a political conflict as private governing organizations can be seen as the means of bypassing the normal democratic process to overcome the conflict. Particularly in the case of the blockchain-based financial ecosystem, any decision made by a governance mechanism could significantly affect community participants' economic stakes, which could make the decision-making process more complicated than the case of the Internet. Thus, we need to pay particular attention to the way the governance mechanism makes a decision and the way to legitimize it for the blockchain-based financial ecosystem.

However, the question is how to actually derive a governance mechanism for the blockchain-based financial ecosystem. In the case of the Internet, non-government stakeholders get together to develop protocols and standards to make the Internet open and globally consistent and have established organizations that are in charge of the day-to-day operation. However, the blockchain community is still fragmented and there seems to be no organic emergence of the cross-project governance mechanism, which is one of the stark differences between the Internet and blockchain ecosystems. Assuming the low probability of organic development of such mechanism within the blockchain-based financial ecosystem, we have to identify similarities and differences between the two systems to trigger its development.

4.4 How Internet Standards Are Created

Technology and operational standards are the basis for making a technology globally operable. At ISO, an international standardization body, the standard is decided by voting among nations. On the other hand, in the case of the Internet, anyone can present a new technological proposal. The Internet Engineering Task Force (IETF) is an organization that publishes standard documents of Internet technologies. Anyone can join the process at the IETF as an individual. The standardization process is conducted over mailing lists, and IETF holds three in-person meetings per year. Proposals are discussed by participants and published after a rough consensus is reached. The principle of “Running Code and Rough Consensus” is a good basis for gathering all stakeholders in one place and collaborating to facilitate permissionless innovation.

4.5 Success and Failure in the History of the Internet

The main success of the process of standardization of the Internet is creating a common standard lead by engineers who create running code. However, in the initial phase of the Internet, users believed that they could trust each other. Thus, the design of security and privacy mechanisms was not well considered at that time. This is one of the reasons why many security incidents and privacy problems occur regarding today’s Internet. In the case of financial applications, such a failure may have a huge impact on regulatory goals. Hence, we need to avoid similar mistakes for financial applications.

5 Designing a New Style: Multi-Stakeholder Governance for Finance

With the lessons from the Internet in mind, we discuss the possible governance structure for blockchain-based financial ecosystems. Our aim is to convince readers that regulators and other stakeholders should work jointly to establish a multi-stakeholder governance mechanism to strike a better balance between fully enjoying the benefits of the innovation and mitigating the risks that it could bring.

5.1 Implications from Lessig's Four Modes

In the previous section, we discussed the fact that regulators face difficulty in achieving regulatory goals using the conventional regulatory approach. What is the alternative approach that regulators can take? Given that the blockchain-based system works Internet and exists within its architectural designs, it is probably wise to adopt a well-known concept of the governance mechanism for the activities in cyberspace, that is, a Lessig's framework known as "pathetic dot theory (Lawrence Lessig 1999)."

Lessig argued that the dot, a target to regulate, is controlled by four enforcement tools: law, norm, market, and architecture. The law and regulation are classic tools for regulators to impose certain constraints on the target, but it is not the only factor affecting people's behavior in cyberspace. Social norms could also be a constraint on the target's behavior through, for example, a tacit understanding in the community of do's and don'ts. In addition, if the behavior is associated with economic motives, market mechanisms such as cost and benefit will also affect the target's behavior. Finally, architecture could also affect the target's behavior by defining technological capabilities and limitations. It should be noted that these four factors are interconnected, and a change in one constraint could change the effectiveness of others.

As such, Lessig's four modes of control can be used as a starting point to consider alternative approaches for regulators. De Filippi and Wright demonstrate what regulators can do to ensure regulability over the blockchain-based network utilizing four different modes of control (De Filippi and Wright 2019). In the following sections, we introduce part of the arguments put forward by De Filippi and Wright.

5.1.1 Laws and Regulations

Above all, De Filippi and Wright suggest that regulators can still utilize law and regulation as an enforcement tool even in a blockchain-based financial ecosystem. For example, regulators can exploit tracing technology to locate a user's address and take enforcement actions directly against end users. In addition, regulators let users be held liable for maintaining the undesirable blockchain-based service. The government may consider punishing them to maintain the network on the grounds that the blockchain network is ultimately maintained by users interacting with the network by, for example, paying fees. This could incentivize people to stay away from services that regulators deem risky or unlawful.

Another target they highlight is the transportation layer, such as ISPs and other information transportation systems. ISPs can detect users who interact with specific blockchain services by utilizing technologies such as deep packets since not all blockchain deploy techniques to hide the data from ISPs. Thus, regulators may regulate ISPs to filter information from an undesirable blockchain-based network. In addition, they also mention information intermediaries as another target because users gain knowledge of the services on the Internet through search engines and social networks.

Moreover, they argue that regulators can affect the behavior of miners/mining pools and other transaction processors within the blockchain network by imposing necessary regulatory actions. Regulators can tweak incentives for miners by, for example, allowing limited liability or providing a safe harbor if miners follow the regulation and process only legal smart contracts. Thus, regulators could increase the power to control the system as such parties play a significant role in operating the system directly. In addition, regulators could have the option to increase or decrease tax depending on miners' behavior.

Finally, they suggest the possibility that regulators impose regulations on those who develop technologies. They can mandate developers to include specific features such as government backdoors within the protocol itself or choose to prosecute developers who intentionally develop undesirable software that helps illicit use. In a similar manner, regulators could also regulate hardware manufacturers' development of devices connected to blockchain networks.

5.1.2 Market Mechanism

Every transaction and smart contract implementation on blockchain can be seen as an economic transaction. Thus, economic incentives can influence the behavior of network participants. Taking advantage of these characteristics, De Filippi and Wright mentioned the possibility that regulators affect the network by intervening in this economic mechanism. For example, regulators can influence the underlying cost structure of a blockchain network by acting as a miner. If regulators take the majority control of the mining power, they can change the required computational power to create incentives/disincentives for other miners to let them abide by regulation. Alternatively, regulators can decide which protocol to support. As another example, regulators can also intervene in the crypto-fiat exchange market because miners decide whether they invest in the hardware to mine the blocks of a

particular blockchain network by comparing the costs of investment denominated in the fiat currency and benefits from the reward and fees of mining denominated in the exchange rate between the fiat currency and crypto. These market interventions could put pressure on the market participants, empowering regulators to change or implement the underlying protocol of the blockchain.

5.1.3 Social Norms

Social norms also influence blockchain-based communities. De Filippi and Wright show, as an example, that each blockchain-based community has a different norm around the hard-fork (De Filippi and Wright 2019). The Bitcoin community values the notion of immutability, while the Ethereum community decided to implement a hard-fork to modify the data on the blockchain as a countermeasure to undesirable incidents such as the Decentralized Autonomous Organization (DAO) case. As this exemplifies, social norms have a profound impact on how community governments could indirectly regulate the operation of the network if they find a way to affect the social norm of the blockchain-based community.

5.1.4 Architecture

In the code-based system, code ultimately determines the course of action that users of the system can take and thus determines the level of control. As opposed to law and regulation, which cannot entirely prevent people from taking specific actions, the code is a rigid rule that strictly limits users' behavior within the system. Thus, De Filippi and Hassan state that "code is increasingly employed in a wide variety of sectors to regulate behaviors, either jointly with, or in addition to, existing laws." They refer to this notion as "Code as Law (De Filippi and Hassan 2016)." In a similar manner, De Filippi and Wright argue that, within the blockchain-based financial ecosystem, regulators can utilize code as an enforcement tool to guide the behavior of the participants therein (De Filippi and Wright 2019).

5.2 How to Fix the Problem

Taking the above analysis into consideration, one possible solution to address regulatory challenges is to encourage the development of code that allows regulators to achieve regulatory goals and to encourage each stakeholder to

use that code. However, in the case of the Internet, it was quite difficult to change the nature of the code, norms, and market to address social problems it caused as all of them had already been fixed to a great extent. Regulators and stakeholders within the ecosystem had to work together to influence the way people use code, which was not an easy task not only for regulators but also the community at large. However, unlike the case of the Internet, most blockchain technologies are still in their very early stage of development and there is currently no well-established technical standard. In addition, it seems that the social norm of each community is not as solid as that of the Internet community. Thus, we believe it is possible to encourage development of the code, norms, and the market in a way that ensures the achievement of regulatory goals. Therefore, regulators should start taking actions promptly to ensure sufficient time and flexibility to influence these factors within the blockchain-based financial ecosystem.

We have discussed that leveraging and aligning Lessig's four factors are essential, and multiple stakeholders should cooperate to achieve this. However, the question is how can they align all four factors? We propose the establishment of a platform as a first step to create a healthy governance mechanism in which all stakeholders come together to discuss how to develop codes, regulations, businesses, and norms. For example, when engineering communities and other stakeholders need to align their interests, regulators are expected to provide their regulatory perspectives and pain points, businesses to provide their business needs and obstacles, users to provide their interests, and civil society to provide concerns related to broader social issues in the engineering community so that they can develop codes with these different perspectives taken into account. On the other hand, the engineering community and other parties could input their perspectives so that regulators develop regulatory policies that make sense to the ecosystem as a whole. What is important in this process is to encourage or incentivize participants to communicate with others to develop mutual understanding to make the activities of the platform more sensible and acceptable in society. If this mechanism works well, regulators may be able to develop appropriate regulatory frameworks such as regulatory safe harbors or sandboxes taking technologies and practices within the ecosystem into consideration. Every participant in the platform should try to establish common ground and shared values, which would eventually generate a trust relationship among them. In addition, the governance mechanism needs to be continuously checked to determine if the four factors are working well because these factors could change its characteristics as technology and society continuously transform. If the stakeholders found any emerging issues, the governance

mechanism should provide an appropriate place in a timely manner to discuss how the ecosystem should apply the four factors to tackle the issue. In some cases, regulation may be the best option to address the issue, and technology may do so in other cases. Another case may necessitate a combination of new technology and regulations to develop a workable solution. This flexible and dynamic process could not be achieved through individual efforts within a single stakeholder. Thus, we advocate establishing a multi-stakeholder platform as a place for communication and cooperation among stakeholders.

As elaborated in the next chapter, we have established a multi-stakeholder platform called the Blockchain Governance Initiative Network (BGIN) to serve as an open, global, and neutral place for dialogues among various stakeholders. Before discussing the details of the BGIN, in the next section, we explain the background theory and key considerations for such a platform.

5.3 Multi-Stakeholder Platform for Healthy Governance Mechanisms

Previous sections demonstrated why regulators need to adopt a multi-stakeholder approach to establish a healthy governance mechanism to achieve their regulatory goals. However, as indicated above, such a governance mechanism is unlikely to emerge organically, at least at this moment in time. Thus, regulators should consider playing an active role in establishing a multi-stakeholder platform. Although it is true that the governance mechanism needs to cope with various problems that will change over time depending on the development of the ecosystem and the surrounding world, we attempted to propose some examples of agendas, which we consider relevant for an illustrative purpose, to better understand the role and function of such a mechanism.

Some issues are interrelated to technology and regulation. For example, the easiest and most obvious one is finding a balance between privacy and traceability of blockchain-based financial activities. As seen in previous chapters, both privacy and traceability have important social values, and each stakeholder will have different views on their value. Therefore, stakeholders participating in the governance mechanism need to discuss and find the optimal level of privacy and traceability and find a practical solution to achieve this. This should not be a one-shot discussion because the optimal level could change over time, and the way to achieve it would also change depending on the available technology. Although this agenda mainly deals

with technology and regulation, the decision regarding the governance mechanisms requires input from other stakeholders such as businesses and users as the effectiveness of the solution is subject to the degree of acceptance by the stakeholders in the ecosystem.

Other issues could be more connected to technology, such as interoperability and security. Even though these issues may primarily be associated with technology itself, all other stakeholders need to contribute to the discussion because developed technology would influence the effectiveness/needs of the regulation, business decision, and usability. As blockchain technology is still in its infancy, technical architecture would be an important topic to be dealt with in the governance mechanism, which would help innovation in technology and business models.

Another key consideration is how to convince each stakeholder to participate in the governance mechanism. We have presented a couple of reasons why regulators may wish to have a multi-stakeholder governance mechanism. However, if they want to establish a workable multi-stakeholder platform, the platform should have a proper incentive mechanism so that stakeholders are willing to join the activities. In particular, part of the engineering community working on the blockchain protocols has a strong mindset of freedom and may be reluctant to have a relationship with the government, which is the so-called cypherpunk philosophy. If regulators prioritize their own interests and pay little attention to the needs and pain points of the engineering community, these cypherpunks will not participate in the activity. To consider the incentive mechanism for the engineering community, it would be helpful to understand how open source software (OSS) communities work. Although it would require a lengthy discussion to cover all OSS, there are existing studies dealing with incentive mechanisms for individual engineers participating in the community and governance mechanisms for them. Moreover, we should pay close attention to the governance structure of the platform. While there are many different blockchain projects such as Bitcoin and Ethereum, it is critical to find a way to legitimize the multi-stakeholder governance mechanism as an ecosystem-wide activity. To legitimate the mechanism, we should pay attention to the saying “God is in the detail,” meaning fairness and openness in the governance structure, such as organizational structure, membership policy, decision-making process, funding scheme, and agenda-setting are very important. In fact, the recommended guiding principles of the Internet governance set by the ISOC includes “Open, inclusive, and transparent participation,” “Consensus-based decision-making,” and “Collective stewardship and empowerment” to guarantee the continuous innovation

and growth of the Internet. This is true of the governance mechanism for the blockchain-based financial ecosystem.

Another consideration should be given to the relationship between newly established governance mechanisms and existing structures such as international regulatory forums and blockchain engineering communities. For example, the international financial regulatory framework is currently developed within international forums, such as the Financial Stability Board and the Basel Committee on Banking Supervision. The question we need to consider is whether the governance mechanism we are discussing should directly deal with regulation. At this moment, we believe it is highly unlikely that regulators decide to delegate all regulation-setting powers to private multi-stakeholder governance mechanisms in the near future, as there is no well-established relationship among stakeholders, at least when the governance mechanism is established. If we try to force them to give up their power to set regulation at the onset, regulators would be reluctant to support the new initiative given the financial system and its regulation have a long history and the existing structure is very sturdy. However, the good news is that the G20 leaders agreed on the importance of the collaboration between regulators and other stakeholders to address key issues in decentralized financial systems, as elaborated in the next chapter. Thus, our idea is that stakeholders come together to discuss issues and develop recommendations that stakeholders return to the existing structure to make a decision on their own. As previously discussed, some of the issues may be handled by a combination of certain regulations and new technology, and others could be addressed by a combination of norms and new business models. Since the outcome is developed in the governance mechanism in which each stakeholder participates, all stakeholders involved are required to take accountability for the decision. If regulators decide to develop regulations that are not in line with the outcome from the governance mechanism, they need to explain the reason well enough to convince other stakeholders and the general public of their decision. The same is true for the engineering community.

As indicated in the previous chapter, simply bringing different stakeholders together in one place would not be enough to make real progress. Instead, it is important to incorporate certain mechanisms that ensure that the outcome from the governance activity has a real impact. In fact, several institutions have already worked on issues related to governance, but many of their outcomes are not referred to by key players in the ecosystem, such as engineers and regulators. Therefore, it is important that wider community participants actually take these materials as guiding principles in day-to-day decision-making to generate real impact.

For example, if the participants really want to have interoperability among different protocols, engineers would choose to get together to discuss the standard and try to agree on a single standard even when they have different opinions and different purposes. On the contrary, if they do not need to have a strong desire to have a single standard, they do not need to make a compromise and eventually walk away from the governance structure. In the case of the Internet, as the critical resource management of the Internet mainly rests on ICANN and related entities, once ICANN decides rules, all network participants utilizing the Internet protocol need to follow the rule. In this way, a successful governance structure in the case of the Internet appears to have an inherent nature that allows participants to abide by the outcome from the governance activities. On the contrary, given that there currently seems to be no strong desire for interoperability among different blockchain protocols and day-to-day operations of the blockchain network is distributed to a large number of miners, it would not be enough to simply copy the governance model of IETF or ICANN to allow governance participants to abide by the outcome. Learning lessons from their experiences, we have to continue to seek our own ways to adequately address important issues in blockchain-based financial ecosystems.

6 The Way Forward

6.1 Discussion at G20 and OECD

In 2019, the G20 financial ministers and central bank governor meetings were held in Fukuoka Japan. At the meetings, the group discussed the importance of multi-stakeholder dialogue and then concluded this idea in the communique as follows.

Technological innovations, including those underlying cryptoassets, can deliver significant benefits to the financial system and the broader economy. While cryptoassets do not pose a threat to global financial stability at this point, we remain vigilant to risks, including those related to consumer and investor protection, anti-money laundering (AML), and countering the financing of terrorism (CFT). We reaffirm our commitment to applying the recently amended FATF Standards to virtual assets and related providers for AML and CFT. We look forward to the adoption of the FATF Interpretive Note and Guidance by the FATF at its plenary later this month. We welcome IOSCO's work on cryptoasset trading platforms related to consumer and investor protection and market integrity. We welcome the FSB's directory of cryptoasset

regulators, and its report on work underway, regulatory approaches, and potential gaps relating to cryptoassets. We ask the FSB and standard setting bodies to monitor risks and consider work on additional multilateral responses as needed. We also welcome the FSB report on decentralized financial technologies and the possible implications for financial stability, regulation, and governance, and how regulators can enhance dialogue with a wider group of stakeholders. We also continue to step up efforts to enhance cyber resilience and welcome progress on the FSB's initiative to identify effective practices for response to and recovery from cyber incidents.

Three days before the meeting at G20, the Financial Stability Board (FSB) published a report on decentralized financial technology. This report considers the financial stability, regulatory, and governance implications of DLT and P2P, sets out the benefits and risks of increased use, and underscores the importance of multi-stakeholder dialogue (Financial Stability Board 2019).

In 2019, OECD formed the Blockchain Expert Policy Advisory Board (BEPAB) to create principles for the blockchain ecosystem. In the principles, the OECD concludes the importance of multi-stakeholder collaboration.

6.2 A New Initiative: BGIN

After the G20 Fukuoka meeting, a group of universities named the BSafe.network organized a series of eight workshops for multi-stakeholder discussion in diversified locations around the world. Regulators, engineers, blockchain business entities, and academia join to discuss issues on privacy, identity, key management, security, and so on. Through these workshops, the importance of multi-stakeholder discussion is confirmed.

On March 10, 2020, a group of 23 blockchain and governance experts initiated a new network called the Blockchain Governance Initiative Network (BGIN). BGIN will play a leading role in developing healthy governance in the blockchain-based ecosystem in a similar way as IETF and ICANN do for the development of Internet governance. To serve this purpose, BGIN tentatively aims at:

1. Creating an open, global, and neutral platform for multi-stakeholder dialogue.
2. Developing a common language and understanding among stakeholders with diverse perspectives.
3. Building academic anchors through continuous provision of trustable documents and codes based on an open source style approach.

BGIN is currently in its early stage and actively seeks contributors to this initiative so that it can accommodate diverse opinions from a wider range of stakeholders. The online/in-person discussion is open to all stakeholders, and anyone can contribute to the documentation process based on the philosophy of “Rough Consensus and Running Code.” The diversified community and the openness are essential to create material outputs, such as technical specifications, which could serve to standardize protocols and design proper layers for the architecture.

References

- De Filippi and Hassan. 2016. Blockchain technology as a regulatory technology: From code is law to law is code. *First Monday* 21 (12).
- De Filippi and Wright. 2019. *Blockchain and the law the rule of code*. Harvard University Press.
- Financial Action Task Force. 2012. The FATF recommendations: International standards on combating money laundering and the financing of terrorism & proliferation [Updated October 2020]. <https://www.fatf-gafi.org/media/fatf/documents/recommendations/pdfs/FATF%20Recommendations%202012.pdf>.
- Financial Action Task Force. 2015, June. Virtual currencies: Guidance for a risk-based approach. <https://www.fatf-gafi.org/media/fatf/documents/reports/Guidance-RBA-Virtual-Currencies.pdf>.
- FINCEN. 2020, March. SAR stats—Depository institutions. <https://www.fincen.gov/reports/sar-stats>.
- Gary Basin. 2018, June 21. The state of decentralized exchange, Hacker Noon. <https://hackernoon.com/the-state-of-decentralized-exchanges-235064446ab0>; Moneropedia, supra note 40; Zcash, supra note 38.
- John Armour, et al. 2016. *Principles of financial regulation*. Oxford: OUP.
- Lawrence Lessig. 1999. *Code and other laws of cyberspace*. Basic Books.
- Lightning. 2019. The technical standard for Lightning Network deploys onion routing to enhance privacy. See BOLT #4: Onion Routing Protocol, Lightning Network. <https://github.com/lightningnetwork/lightning-rfc/blob/master/04-onion-routing.md>. Last visited November 1, 2019.
- Yuta Takanashi, et al. 2019. Call for multi-stakeholder communication to establish a governance mechanism for the emerging blockchain-based financial ecosystem, Part 1 of 2. *Stanford Journal of Blockchain Law & Policy*. <https://stanford-jblp.pubpub.org/pub/multistakeholder-comm-governance/release/1>.
- Zcash. 2019. What are zk-SNARKs? <https://z.cash/technology/zksnarks.html>. Last visited November 1, 2019.



Law and Blockchains

Stephen McKeon and Derek Edward Schloss

1 Introduction

The emergence of blockchain-based assets and systems of record have spurred numerous, and sometimes differing, interpretations within segments of the legal code. The decentralized nature of the technology is not always congruent with existing laws and precedent, which typically contemplates an environment with a higher degree of centralized control. In this chapter, we review various aspects of the legal environment and review how applications such as smart contracts, cryptocurrencies, tokenized securities, and decentralized autonomous organizations (DAOs) are treated with regard to tax law, intellectual property (IP) law, and securities law. We focus primarily on U.S. law, but comment on the global legal environment where applicable.

We begin with securities law because it is often unclear whether network-based assets constitute financial securities or not. We review the Howey Test and point to regulatory exemptions commonly utilized by issuers. Additionally, we briefly review the proposed benefits of issuing securities on-chain.

S. McKeon (✉)

Lundquist College of Business, University of Oregon, Eugene, OR, USA

e-mail: smckeon@uoregon.edu

D. E. Schloss

Collab+Currency, Eugene, OR, USA

Smart contracts represent another area of intersection with the legal environment, specifically with contract law. A smart contract is a software code that has dominion over the value to be exchanged and executes an outcome autonomously based on a set of pre-specified conditions. They promise to reduce ambiguity, and increase speed and efficiency of the contracting environment. However, whether or not these agreements are deemed to be legally enforceable contracts has been the source of debate.

Ultimately, decentralization creates legal challenges. For example, GDPR did not contemplate data on distributed ledgers, bringing privacy questions to the fore. Illegal trade is facilitated by pseudonymous money that can be transmitted at distance without an intermediary. And the rise of decentralized autonomous organizations, where governance is executed through code, raises a host of new legal questions.

The rest of the chapter is organized as follows: Sect. 2 covers securities law, Sect. 3 examines the literature around smart contract law, Sect. 4 examines the unique challenges around decentralization, Sect. 5 reviews tax law around cryptocurrencies, Sect. 6 reviews the intersection of blockchains and intellectual property law, and Sect. 7 concludes.

2 Securities Law

The application of securities laws to blockchain tokens is the subject of considerable debate and an active segment of literature. A token is simply a digital wrapper that allows ownership of the asset to be recorded on a distributed ledger. The challenge is that this functional form can endow the holder with any number of rights, therefore, the question of whether a particular token is subject to securities laws often boils down to the specific rights associated with ownership. In this section, we review the foundational framework for securities law in the United States, synthesize literature examining tokenization and securities laws, and cover purported benefits of representing financial securities on-chain as a direction for future research.

We begin by outlining what constitutes a legal security in the United States. After the U.S. equity markets crashed in 1929, Congress began work on a regulatory cleanup that aimed to prevent some of the core problems that occurred during the lead-up to the crash—specifically, speculation and information asymmetry between issuers and investors.

As a result, Congress created the Securities Act of 1933, which defined a number of financial instruments as securities, including stocks, bonds, notes, security futures, participation in profit-sharing agreements, and investment

contracts (15 U.S. Code § 77a The Securities Act of 1933). For the first time, the Securities Act of 1933 provided clarity as to the federal classification of securities, and when securities laws would apply to the sales of these instruments (15 U.S. Code § 77a The Securities Act of 1933). Many of the instruments listed in the Securities Act of 1933 are fairly straightforward, and include clearly defined instruments like stocks and bonds. However, one area that has seen enormous legal discussion, especially as applied to blockchain-based instruments, is the investment contract.

In 1946, the U.S. Supreme Court clarified the classification of an investment contract with a four-prong analysis called the Howey Test, which continues to be utilized today (Securities and Exchange Commission v. W. J. Howey Co., 328 U.S. 293). Under the Howey Test, an investment contract is any contract, transaction, or scheme, whereby a person (1) invests money, (2) in a common enterprise, (3) and is led to reasonably expect profits, (4) derived from the efforts of others (Securities and Exchange Commission v. W. J. Howey Co., 328 U.S. 293).

In April 2019, the SEC offered guidance on the application of the Howey Test to digital assets when it published its Framework for “Investment Contract” Analysis of Digital Assets.¹ In writing the framework, the SEC wanted organizations considering the “offer, sale, or distribution” of a digital asset to apply the Howey Test framework and the SEC’s digital asset framework to determine whether securities laws should apply.

However, operationalizing the Howey Test is not a straightforward exercise because tokenized networks contain features not previously contemplated in investment contracts and do not map cleanly to historical precedent and case law. For example, many tokens only endow the holder with the ability to access digital goods or services from the network, and don’t contain any cash flow or governance rights. Given the large variation in design of assets on distributed ledgers, several efforts have been made to provide a framework or taxonomy for categorization, notably Rauchs et al. (2018). Additionally, Henderson and Raskin (2018) offer suggestions on how to operationalize Howey in the context of blockchain networks. One idea they advance is the creation of a “Bahamas Test” to determine the degree of decentralization. This builds on the notion that it is possible for an asset to be a *transitory* security; regulated as a security at inception, but subsequently evolve to a degree of decentralization as the network matures to where it no longer constitutes a financial security in the eyes of the SEC and regulatory jurisdiction passes to another body, such as the CFTC.

¹ Available at <https://www.sec.gov/files/dlt-framework.pdf>.

Drew Hinkes, an attorney focusing on blockchain, proposes four categories of security token offerings.² First, security-wrapped ICOs (SICOs) are network assets or utility tokens offered pursuant to registration exemptions so that their offering complies with U.S. or other securities frameworks. Second, tokenized equity or debt securities (TEDs) are traditional securities, like equity or debt securities, issued in digital token form. Third, tokenized asset-backed securities (TABs) are digital tokens that represent an ownership claim against, or ownership share in, an asset such as gold or pool of assets such as a venture capital fund. Finally, Transactional Security Instruments (TSIs) are securities that could be redeemed directly by the issuer for goods and services.

Decentralized network assets are often sold through an Initial Coin Offering (ICO) and existing regulatory approaches to ICOs are covered in Chapter 20 of this handbook, as well as Gurrea-Martínez and Remolina (2018), Park (2018), and Maas (2019). While the debate over applicable law for decentralized networks will likely persist for some time, we note that the legal environment is clearer for the second and third categories in Hinkes's taxonomy, where the token is simply a new digital representation of financial securities with which we are already familiar.

In the United States today, securities offerings made to U.S. residents must either be registered with the SEC or exempt under the Securities Act of 1933. Registration with the SEC means the issuer creates a public offering for the security, and the issuer can raise money from both accredited and unaccredited investors. Public registration also carries a number of other benefits, including the ability to publicly advertise the offering, along with immediate trading and liquidity of the registered securities.

Alternatively, issuers can also raise money through a private placement by filing for an exemption under the Securities Act of 1933. Before the JOBS Act was signed into law in 2012 by then-U.S. President Barack Obama, private placement issuers could only raise money from accredited investors. As it relates to natural persons, accredited investors are defined under U.S. securities laws as any natural person whose individual net worth, or joint net worth with that person's spouse, exceeds \$1.0M (17 CFR § 230.501). The accredited investor label can also be satisfied by any individual who generated \$200k in each of the two most recent years, or enjoys joint income with that person's spouse in excess of \$300k in each of the two most recent years, and has a reasonable expectation of reaching that same income level in the current year.

² See <https://www.coindesk.com/the-security-token-market-needs-better-lingo>.

Both the private placement rules and accreditation standards were largely constructed to protect unsophisticated individuals from predatory investment opportunities. In 2012, the JOBS Act created the ability for non-accredited “retail” investors to participate in certain exempt private offerings under Reg D, Reg A+, and Reg CF. While these private exemptions have created new wealth-generating opportunities for nonaccredited investors, in June 2019, the SEC requested public comment on ways to “simplify, harmonize, and improve” the private offering securities exemption framework in the United States (“SEC Seeks Public Comment on Ways to Harmonize Private Securities Offering Exemptions,” <https://www.sec.gov/news/press-release/2019-97?hootPostID=7526bcf08a34ed3da2c1f335c8e39d1f>).

Today, the most common types of private placement exemptions used in digital asset issuances, along with a brief description of their features, are listed below. A more thorough discussion of the exemptions used in tokenized security offerings can be found in Goforth (2019).

Regulation D Rule 506(c) is a registration exemption that provides for an unlimited capital raise, the ability to solicit, along with an exemption from the state “blue sky” securities laws (17 CFR § 230.506(c)). However, the use of this exemption limits the pool of potential investors to accredited investors only, and the issuer must take reasonable steps to verify each investor’s accreditation status. In addition, securities offered and sold under Reg D Rule 506(c) cannot be transferred for twelve months, with limited exceptions.

Regulation D Rule 506(b) has similar features to Reg D Rule 506(c) (i.e. unlimited capital raise, twelve-month transfer restrictions), except no general solicitation is allowed to market or advertise the offering, and no accredited investor verification is required, as long as the company has no reason to believe that any investor is not an accredited investor (17 CFR § 230.506(b)). However, Reg D Rule 506(b) does allow up to 35 unaccredited investors to participate in the offering.

Regulation A+, often referred to as the “mini IPO,” allows issuers to offer and sell securities pursuant to general solicitation, with no minimum investment, that will be immediately transferable once delivered to the investor (17 CFR § 230.251). Here, securities can be offered and sold to both accredited and unaccredited investors, and the issuer’s capital raise can be as much as \$20.0M under Reg A + Tier 1 and \$50.0M under Reg A + Tier 2. However, drawbacks include a lengthy SEC review and process, audited financial statements in the offering statement, and ongoing annual financial reporting. The first two token-based Reg A + s were approved in July 2019: Blockstack and Props.

Regulation CF, also known as the “crowdfunding exemption,” also allows for issuers to offer and sell securities to both accredited and unaccredited investors (17 CFR § 227.100). However, while Reg CF does feature the ability to raise from a diverse investor pool, the issuer’s capital raise is capped at \$1.07M, securities transfers are restricted for twelve months, and issuers face reporting obligations.

As described above, there are numerous regulatory paths to issue securities on blockchains. A related question is why an issuer would wish to do so. The answer is that utilizing a digital wrapper that allows the asset to trade “on-chain” enables various features that we don’t observe in securities today, and raises numerous questions about how securities will evolve. We summarize these briefly below and a more thorough review can be found in McKeon (2018).

24/7 Markets: Today, the major U.S. stock market exchanges open at 9:30 a.m. and close at 4:00 p.m. (EST) on weekdays. Electronic communication networks (ECNs) allow expanded trading hours, but are not accessible by most retail investors. However, the vast majority of trading venues for digital assets like bitcoin operate 24 hours per day, seven days a week. As traditional financial assets like stock and bonds begin to be issued in the form of tokens, a question is whether around-the-clock trading will be the norm, and further, whether that is an optimal outcome. Barclay and Hendershott (2004) report that stock prices after-hours are less efficient than prices during the day and are characterized by large bid-ask spreads. However, Barclay and Hendershott (2003) find that the low trading volume observed during off-hours can facilitate price discovery.

Rapid Settlement: Exchanges like NASDAQ and NYSE can execute trades very quickly, but settling these transfers takes time. In 2017, the SEC adopted a shortened settlement cycle for most broker–dealer transactions to T + 2.³ Settling transfers of private securities, like LP and LLC interests, can take even longer. When ownership claims are tokenized on a distributed ledger, settlement can occur nearly instantaneously. Standardizing settlement has numerous implications, one of which is cross-border flows. Bekaert (1995) finds that inefficient settlement systems are a friction that creates an indirect barrier to investment, suggesting that tokenization may enhance cross-border investment flows.

Cost Reduction: Ritter (1987) outlines two categories of costs when firms go public—direct costs such as legal and underwriting, and indirect costs such as underpricing. A commonly cited benefit of security tokens are reductions in

³ See <https://www.sec.gov/news/press-release/2017-68-0>.

back-office costs associated in the issuance of securities through automation. Additionally, we may see higher utilization of auction processes instead of traditional book building, potentially reducing both underwriting fees and underpricing.

Servicing: After issuance there are ongoing costs around servicing securities that may be reduced through automation. For example, when startups are acquired, reconciling the capitalization table to the underlying purchase agreements and option grants is costly. This problem is exacerbated as companies grow. In 2015, a court ruling required Dole to pay all shareholders, and while 36M shares were outstanding, claims for payments exceeded over 45M shares (Solomon 2017). When ownership claims are tokenized, cap tables can be reconciled in real time by code. In addition, contractual features like liquidation preferences and drag-along rights can be programmed into the security token, allowing managers to more easily run scenario analysis to calculate payoffs under different assumptions.

Looking forward: Eventually, digitally wrapped securities will allow us to build in contractual features that have previously been impossible or costly to execute manually, but become feasible through automated referencing. For example, features that tie voting rights to the duration of ownership could be useful in shaping corporate governance and mitigating managerial myopia. Smart securities will also facilitate bundling additional rights, such as early access rights to products or services for investors.

Further, they will allow unbundling of rights. Voting rights could be sold to activists while retaining cash flow rights. Dividend rights could be unbundled like Treasury STRIPS. And companies will unbundle specific revenue streams and finance them independently. Finally, we will see automated referencing between different layers in the capital stack as complex revenue sharing and payment waterfall agreements become much simpler in this environment (Lippiatt and Oved 2018).

Many of these features are dependent on the ability of computer systems and software to exchange and make use of information (i.e. interoperability). The Internet is a stack of protocols that standardize information (i.e. TCP/IP, SMTP, FTP, SSH, HTTP), however, the systems that regulate the transfer of value in our financial infrastructure lack compatibility. The great promise of applying blockchain protocols to securities and financial infrastructure generally, is that they will impose a set of standards that facilitate greater interoperability across asset classes, across borders, and across investor types.

3 (Smart) Contract Law

A contract is an enforceable agreement that, when violated, allows an injured party to access legal remedies. Building on this, traditional contract law is a remedial institution. Its purpose is not to ensure performance *ex ante*, but to resolve the wrongs that might arise *ex post* (Werbach and Cornell 2017).

Traditionally, it's been assumed that contractual agreements must always require the backing of a legal system. However, recent developments in technology have led some to speculate that smart contracts could one day displace contract law (Tapscott and Tapscott 2016), while others argue that smart contracts have little to do with legal contracts (Werbach and Cornell 2017), and/or represent an alternative to the legal system (Savelyev 2017). In this section, we examine the intersection of smart contracts and contract law.

A contract's terms and conditions are interpreted by each party to the contract, however, if a disagreement arises, third parties can be utilized to interpret and enforce the contract's terms. For example, a judge can interpret a disputed contract's conditions, and a local sheriff can enforce the judge's interpretation. Smart contracts are different because execution is automated. The smart contract itself has dominion and control over the physical or digital objects needed to effect execution (Raskin 2017).

The term "smart contract" was coined by Szabo (1996) who defines them as "a set of promises, specified in digital form, including protocols within which the parties perform on these promises." He points to a vending machine as a legacy example of a smart contract, in that it has control over the objects in the transaction and can effect the transaction automatically by issuing an unopened drink, so long as money is inserted into the machine.

Recently, automatic execution is often ensured through a computer running code that has translated legal prose into an executable program. A more advanced example might be a car that has a program installed to prevent ignition if the terms of a debt contract are not satisfied (Raskin 2017).

Although smart contracts were functioning prior to blockchains, they have gained substantial awareness in recent years as blockchains such as Ethereum have emerged as "smart contract platforms." There are three features of blockchains that have led to this rise in smart contracts: (i) acting as a source of truth, (ii) dominion over payment mechanisms, and (iii) ease of deployment through interoperability.

"A blockchain is a general-purpose technology for trusted transactions" (Werbach and Cornell 2017). Blockchains aim to solve the problem of establishing a consensus of information without relying on a single party. How might blockchains and smart contracts work together? The terms of a

contract, and the state of facts relating to the performance of a contract, can be programmed as information within a blockchain to leverage these same truth-verifying benefits (Raskin 2017).

Automation ensures performance, for better or worse, by removing human discretion from contract execution. In addition, blockchains can trustlessly verify that contract's execution. For example, if thousands of nodes on a decentralized network verify that Derek paid Steve \$50 on August 13th at 3:30 PM, one can assume this occurred with a high degree of certainty (Raskin 2017). This is the power of smart contracts and blockchains working in concert, and these smart contracts can be executed with astonishing speed (Wright and De Filippi 2015).

Szabo (1997) points to the cost of breach as a defining feature of smart contracts. Strong smart contracts have prohibitive costs of revocation and modification, while weak smart contracts do not. This extends not only to the parties engaged in the transaction, but to third parties as well. For example, if a court is able to alter a smart contract after it has been executed with relative ease, then it's a weak smart contract. But if the cost of modifying the contract is so high so as to prevent third-party interference, then it's considered a strong smart contract (Raskin 2017). The cost of breach and modification is an important consideration because smart contracts lose some of their core benefits and efficiency if they can be changed easily by an outside third party.

With traditional contracts, executing a contract can be a rigorous process involving numerous intermediaries, which increases the inefficiency of execution. For example, when you buy or sell real property, you might have brokers, attorneys, deed companies, and lenders. In these types of scenarios, the interests of all involved parties must be efficiently organized through contractual agreements.

Smart contracts aim to eliminate some of these inefficiencies by embedding contracted conditions into code. Algorithmic enforcement of a smart contract allows conditions to be executed as quickly and cheaply as other computer code (Werbach and Cornell 2017). Cost savings occur at each stage, from negotiation to enforcement, and potentially replace judicial enforcement with an automated mechanism (Werbach and Cornell 2017).

In the context of a debt instrument like a bond, a smart contract can automate the payment of interest to each investor on pre-specified due dates. In addition, the bond can service itself automatically when triggered by the borrower sending funds to the smart contract. Smart contracts give both the promisor and promisee the ability to encode finality so that parties can organize their behaviors around a bargained-for certainty.

In traditional contracts, ambiguity causes problems. This is observed as far back as *Raffles v. Wichelhaus*, where a controversy arose over a cotton shipment contract because two different ships named “Peerless” were sailing similar routes but at different times of the year (*Raffles v. Wichelhaus*, 2 H. & C. 906 (1864)). The contract was deemed unenforceable due to ambiguity. Unlike traditional contract formation, smart contracts that utilize blockchains have the potential to mitigate misunderstandings and mistakes over the terms of an agreement because they explicitly reference a single source of truth.

Although smart contracts enable some desirable features, their enforceability as legal contracts is in question. A variety of challenges exist that we highlight below.

The intent that matters is objective, not subjective, as intent is manifested by the actions of the parties. Thus, a legal contract will exist for a smart contract only if the actions of the parties, judged objectively, manifest an intention that the smart contract is to be legally enforceable (Werbach and Cornell 2017).

Under common law, to objectively demonstrate that a contract has been formed, there must be offer, acceptance, and bargained-for consideration. To constitute a legal offer, smart contract code could simply be posted to a blockchain’s ledger, verifiable by any downstream party who attempts to engage with the smart contract’s code. To evidence counterparty acceptance of the smart contract, action must be taken to initiate acceptance of that smart contract, such as ceding control over a certain amount of money to the code (Raskin 2017). One question that remains unsettled is whether follow-on contracts that are established autonomously by the first contract are enforceable, as it is not clear that legal intent can be presumed (Giancaspro 2017).

The third element required to constitute a legally enforceable contract is consideration. Courts believe that the mutuality of obligation by both parties, or consideration, distinguishes contracts from gifts, for which parties do not have the same rights of legal enforcement (Raskin p. 322). In the case of smart contracts, consideration can be presented unilaterally, like a vending machine, or bargained-for as in the terms of a loan agreement (Raskin, p. 323).

Once a legally enforceable contract has been formed, a common law contract can either be performed, modified, or breached by a party to the contract. Because smart contracts have the ability to automate execution of agreed-upon conditions, some might argue that the risk of smart contract breach is significantly mitigated. In this way, performance of a smart contract’s conditionals can be relied upon with greater certainty than traditional contracts due to its automated status.

However, one potential problem that could arise is imperfect, but substantial, performance of a smart contract. In the United States, the common law doctrine of substantial performance permits a contract to be recognized even if the performance does not fully comply with the express terms laid out. (RESTATEMENT (SECOND) OF CONTRACTS 237 cmt. D (1981)). However, execution within smart contracts is often binary and incompatible with partial performance where the outcome was not contemplated and specified by the parties during formation of the smart contract (Raskin 2017).

Assuming a risk of imperfect performance exists, what can parties to a smart contract do in order to fall inside the lines of a legally enforceable contract? One way is by baking in a certain degree of discretion into the terms of the smart contract, or by simply not using a smart contract if discretion is required. Alternatively, the smart contract could be structured to permit arbitration. However, incorporating flexible features into the smart contract will inevitably take away from the decentralization and efficiency that make smart contracts an attractive option to begin with (Werbach and Cornell 2017).

Another potential challenge to smart contract performance is modification of the contract before full performance occurs. Under common law, parties to a contract can modify an otherwise enforceable contract's terms, or be absolved from performance entirely, under certain scenarios. The common law doctrines of impossibility and impracticability are two examples of legal excuse from performance. For example, if a contract becomes illegal after it is formed, parties are typically excused from performance with no remedy for aggrieved parties. With smart contracts, however, automation will typically frustrate any potential for post-formation modification (Raskin 2017).

Another potential issue that might accompany smart contracts is the misalignment of contract performance and contract intent. What happens when the outcomes of a smart contract diverge from the outcomes that contract law demands? In the United States and other common law systems, *ex post* enforcement is the preferred system of enforcement. An *ex post* enforcement analysis of smart contract breach might be complicated by examples where smart contracts are followed by the letter of the code, but not necessarily by the spirit of the code.

Smart contracts often operate in tandem. For example, a Decentralized Autonomous Organization (DAO), is a type of firm that utilizes smart contracts to execute its corporate governance. DAOs are discussed in more detail in Section IV.c. of this chapter, but we note here that in DAOs, the rules of management are typically predetermined. In 2016, a smart contract called "The DAO" was formed on Ethereum, a public blockchain and funded with \$150 million by numerous participants. The intention was to function

as a venture capital fund. However, shortly after the money was raised, \$40 million was siphoned from the pool by a hacker who used the smart contract in an unanticipated way.

In the case of The DAO, the thief's engagement with the smart contract was orthogonal to the intent of the contract. If legal remedies were pursued, the common law requirement that each party demonstrate an "objective intent" to enter into an agreement would likely not be satisfied, as the terms being agreed upon by the thief were materially different than the terms being offered by The DAO. As a result, The DAO's smart contract would fall outside of contract law, and The DAO would have no access to legal remedies through a breach of contract argument.

A related issue highlighted by The DAO is that the parties were anonymous. This raises the question as to whether they had the legal capacity to enter a contract. For example, in the United States and numerous other countries, individuals under the age of 18 lack the legal capacity to enter contracts (Giancaspro 2017). This suggests some form of digital ID validating capacity will need to be tied to smart contracts in order for them to be legally enforceable.

Smart contracts are just one piece of a larger trend of technology disrupting human engagement and decision-making. However, as might be clear by the summary above, the introduction of automation into historically judgment-laden fields will create challenges for legal and practical accountability. These challenges notwithstanding, contract law is resilient, and it may evolve as a result of these new technological challenges (Werbach and Cornell 2017).

4 Decentralization Poses Unique Legal Challenges

Although there is little new legislation that addresses distributed ledgers directly, their use raises a variety of challenges and potential sources of liability under existing laws (Zetsche et al. 2018). For example, Walch (2019) examines whether software developers should be treated as fiduciaries. In this section, we review two areas where substantial questions remain unanswered: data privacy and decentralized autonomous organizations.

4.1 Blockchains and Laws Around Data Privacy

As the world becomes increasingly awash in data, a burgeoning segment of law addresses data privacy. Of particular note is the European Union's

General Data Protection Regulation (GDPR). Finck (2018) addresses how blockchains present a challenge to GDPR as they are currently constructed, and what solutions to the challenge might look like moving forward.

GDPR contains four principles with respect to personal data:

1. Data minimization
2. The right to amendment
3. The right to access
4. The right to be forgotten

Each of these principles is challenging in the context of blockchains in part because resistance to censorship and immutability of data is a defining feature of blockchains.

The GDPR requires full anonymization before data is no longer considered to be personal data. With regard to blockchains, Finck notes that this means that pseudonymous data, even if it is encrypted or subject to a hashing process, continues to be classified as personal data and therefore subject to GDPR.

If data on blockchains is personal data and therefore subject to GDPR, then data controllers are compelled to comply. In centralized systems, which is what the authors of GDPR had in mind, data controllers are easy to identify and hold accountable. Examples include companies such as Facebook or Google.

However, in blockchain networks, there is no centralized data controller. Rather, data is processed by all nodes on the network, which can number in the thousands. Further, even if all the owners of nodes could be identified (a near impossibility for large networks), they are typically spread out across the globe. Processing of personal data in foreign countries is potentially a further violation of GDPR, however, enforcing compliance on a decentralized network of nodes is infeasible. Finck therefore concludes that many blockchain networks that exist today are fundamentally at odds with GDPR.

Although blockchains and GDPR appear to be incompatible, Finck points to a number of possible solutions. For example, courts could declare that hashed data is adequately anonymized. However, the more likely near term solutions are technical in nature. Cryptocurrencies focused on privacy such as Z-cash and Monero have pioneered application of zero-knowledge proofs and ring signatures, respectively. Zero-knowledge proofs allow verification of transactions without revealing details of the transaction. Ring signatures obfuscate transaction details by tying multiple keys together making it impossible to determine which one was used by the sender. In sum, blockchains

have great promise for portability and although they are currently incompatible with privacy, there is hope that this challenge is surmountable in time.

4.2 Illegal Trade

The pseudonymous feature of Bitcoin and other cryptocurrencies gives rise to their use in illegal transactions. While most of this chapter focuses on legal interpretation of blockchains, we briefly review the literature on the use of blockchain-based assets for illegal trade.

Perhaps the first large-scale use case for Bitcoin was as a payment method on Silk Road, a marketplace for black market goods. Prior to 2013, when the U.S. government shut down Silk Road, some (disputed) reports suggest it accounted for up to half of all Bitcoin transactions (Yermack 2017).

More recently, as Bitcoin has gained mainstream adoption, the proportion of illegal trade has declined. Moreover, Rogoff (2016) and Harvey (2014) both point out that cash, particularly the US \$100 bill, is used much more commonly in illegal trade and Harvey suggests that the idea that Bitcoin is mainly used for criminal activity is a myth.

One of the obvious differences between cash and cryptocurrencies is that cash requires physical delivery while cryptocurrency can be transacted at distance. Most regulated forms of electronic cash such as PayPal require real word identification and are therefore less than ideal for illegal transactions. Marketplaces for contraband exist on the darknet, which is beyond the scope of this chapter, but we point the reader to Van Slobbe (2016) for a more detailed description.

Foley et al. (2019) offer a comprehensive account of the methods and scale of cryptocurrency use in illicit trade. The authors use two methods, network cluster analysis (SLM) and detection controlled estimation (DCE), to classify bitcoin users into primarily legal/illegal activity. Using these classifications, they report that illicit users account for about one-quarter of total users and account for approximately 20% of the U.S. dollar volume of transactions. The figures reported in Foley et al. (2019) are substantially higher than those found in other studies such as Soska and Christin (2015) and Meiklejohn et al. (2013).

Bad actors engaging in illegal trade and money laundering using bitcoin impose a negative externality on law abiding members of the cryptocurrency community. For example, obtaining a banking relationship is notoriously

difficult for many firms in the ecosystem.⁴ One of the primary concerns for banks is identifying the source of funds for the depositor, which can prove challenging for firms that trade in Bitcoin, for example, a spot exchange. However, the transparency offered by the Bitcoin ledger allows some analysis to be done on the wallets through which the coins have passed. Firms such as Chainalysis, Elliptic, and Elementus offer products to financial institutions that look back through the ledger history of coins that are deposited to create a risk assessment of nefarious activities. We note that this type of analysis is impossible with cash, and illustrates why distributed transaction ledgers may ultimately be one of the most potent tools available to regulators.

4.3 Legal Considerations Around Decentralized Autonomous Organizations (DAOs)

It is important to recognize that most organizations are simply legal fictions which serve as a nexus for a set of contracting relationships among individuals.—Jensen and Meckling (1976)

One of the more fascinating concepts that blockchain enables is that of a decentralized autonomous organization (DAO). A DAO is an organization where the firm's resources are controlled through the use of one or more smart contracts. Ownership, actions, and value flows in a DAO are strictly dictated by the terms set forth in code. A DAO is a nexus of smart contracts.

This organizational form raises a number of legal questions, for example: jurisdiction. If the organization lives entirely in the cloud, where all interactions are executed peer-to-peer by software, it is not clear what set of laws may apply. Many observers point to the location of the owners, which means many different sets of laws may apply simultaneously, but also gives rise to the problem that identity may be shielded, so enforcement becomes very difficult.

The most famous DAO was simply called “The DAO.” It was to function as a decentralized venture capital fund, where investors could contribute value in the form of ether to a pool, nominate recipients, and vote on the deployment of funds. The pool grew to over USD\$150 million in less than a month. Before The DAO became operational, a bug in the software allowed hackers to drain a meaningful portion of the fund, resulting in a hard fork to roll back the transactions and recover the assets. Detailed discussion of the

⁴ See <https://www.wsj.com/articles/lack-of-banking-options-a-big-problem-for-crypto-businesses-11558092600>.

fork is outside the scope of this chapter, but the relevant fact is that The DAO never went live. From a legal perspective, the SEC later went on to advise that if it had gone live, The DAO tokens would have constituted securities under U.S. law. The question is: upon whom can the SEC impose enforcement? The DAO itself is simply computer code, with no physical address and it did not register with regulators in any nation. This raises the question of how U.S. courts would view the legal organization form of a DAO.

Metjahic (2017) holds that the most appropriate legal form that would have been applied to The DAO is a general partnership, as defined by the Uniform Partnership Act (UPA). When the investors purchased the tokens, in exchange for ether, the tokens granted them proportional cash flow and voting rights, indicating the intention to carry on a business for profit. The transparency of the smart contracts suggests that the investors knew or should have known that they shared a common interest with other investors. Of note is the fact that if all token holders are general partners, they do not have the limitation in liability afforded by other forms of legal organization.

If The DAO was not deemed to be a general partnership, Metjahic (2017) points to a Joint Venture as the next most likely determination. Joint ventures share many attributes with partnerships, but are often interpreted as more limited in duration and purpose. As with partnerships, the existence of a formal agreement is not required as one of the tests that determines this legal status.

Ultimately, DAOs fit somewhere “between an informal online group and a more formalized corporate entity” (Wright and De Filippi 2015, p. 32). Membership may be as fluid as an online group of content creators, and DAO participants will likely include machines in addition to humans. These new dynamics raise new legal questions and may generate a new body of law specific to this new ecosystem. Just as the customs and practices of merchants in Europe developed into *Lex Mercatoria* several centuries ago, and ground rules established at the advent of the Internet formed the basis of *Lex Informatica*, Wright and De Filippi (2015) posit that we are at the dawn of *Lex Cryptographia*.

5 Tax Law and Cryptocurrencies

How cryptocurrencies should be taxed has generated substantial debate and uncertainty for market participants (Lerer 2019). In this section, we review relevant literature on blockchains and tax law.

5.1 Bitcoin as Property

In 2014, the Internal Revenue Service (IRS) in the United States issued Notice 2014–21 providing guidance that virtual currencies should be treated as property, not currency, for the purposes of tax reporting.⁵ This treatment has important implications for economic transactions. Weekley (2018) states “The Commissioner’s Treatment of Bitcoin as property will discourage people from using them in a trade or business.” The friction most often cited is reporting of capital gains and losses. When bitcoin is used to purchase goods and services, the transaction is effectively treated as selling property. Since the exchange rate between fiat currencies like the U.S. Dollar and virtual currencies, like bitcoin, are constantly fluctuating, it means that the price at which one acquired the virtual currency is almost certain to be different than the prevailing market price at the moment of a subsequent transaction. The implication is that purchasing items as trivial as a cup of coffee will trigger a taxable gain or loss. To complicate the matter further, bitcoins are divisible to eight decimal places, so a transaction could contain fractions of many different bitcoins, each with a different basis.

One solution is a *de minimis* exception, which relieves the taxpayer from reporting gains on small transactions. This exception already exists in the U.S. tax code for small gains on foreign fiat currencies. In 2017, Rep. David Schweikert introduced H.R. 3708: “To amend the Internal Revenue Code of 1986 to exclude from gross income *de minimis* gains from certain sales or exchanges of virtual currency, and for other purposes.” It was referred to the House Ways and Means Committee, but has not been voted on as of this writing. An alternative solution is proposed by Weekley (2018), who suggests cryptocurrencies should be treated like frequent flyer miles.

Finally, while the *de minimis* exemption addresses taxation issues around the purchase of goods and services, a related but separate tax issue arises when one cryptocurrency is exchanged for another, for example, trading bitcoin for ether. Currently, this triggers a gain or loss, which is taxable. However, some observers suggest it should be treated similarly to when an investor sells a piece of real estate and rolls the capital into another real estate investment, that is, it should be treated as a 1031 like-kind exchange. In December 2018, Rep. Ted Budd introduced H.R. 7361 (Virtual Value Tax Fix of 2018). As per this bill, gains on crypto to crypto exchanges can be tax deferred as per the Section 1031 of the “Internal Revenue Code Tax Cuts and Jobs Act.” Since a

⁵ IRS guidance: <https://www.irs.gov/pub/irs-drop/n-14-21.pdf>.

token can represent any asset, a challenge that remains with this approach is determining with more specificity what assets would be eligible.

5.2 Hard Forks

One of the most interesting aspects of cryptocurrency tax law from an academic standpoint is the treatment of “hard forks.” A hard fork refers to instances where a group of developers creates a new version of a virtual currency by making a change to an existing protocol. The result is that the owner now has the asset on the original chain, as well as rights to assets on the new chain. These events are also known as chain splits. An example is that of Bitcoin and Bitcoin Cash. A segment of the Bitcoin community felt that certain changes were needed to the Bitcoin protocol, but they were unable to rally enough support to make the changes. Therefore, they created a new asset, called “Bitcoin Cash,” through a hard fork. A user’s wallet that held 3.5 bitcoin prior to the fork still has 3.5 bitcoin after the fork, but now also contains 3.5 bitcoin cash.

Many readers will immediately recognize the similarities to corporate spin-offs, which are not taxable events, however, there are some important differences. First, a fork requires no consent from participants on the original blockchain or the taxpayer, it is created independently and unilaterally by a third party. That said, relevant case law around “treasure troves,” which are sudden and unexpected windfalls for which the taxpayer bears no responsibility, suggests that taxpayers could be responsible for income taxes on forks even if their receipt is completely passive. Second, a public market may not exist for the new asset on the day it is created—an exchange must be convinced to list it—and therefore there is no price discovery to use for tax reporting.

Several recent papers examine the tax treatment of hard forks including Landoni and Pieters (2019), Webb (2018), and Xu (2018). Webb points to the 1955 Supreme Court case *Commissioner v. Glenshaw Glass Co.* where the court defined income to include:

1. instances of undeniable accessions to wealth,
2. clearly realized,
3. over which taxpayers have complete dominion.

Glenshaw does not provide clarity on cryptocurrency hard forks, because each of these three points is in question. Extreme price volatility and illiquidity may impact the degree to which accession to wealth is undeniable. Further,

clear realization and dominion are affected by where the user stores the assets. For example, many exchanges did not support the Bitcoin Cash hard fork, so there was no way for their clients to claim the new assets even if they held bitcoin in their account at the time of the fork.

Xu (2018) points toward some proposals on how to treat hard forks. The American Bar Association (ABA) submitted a comment letter suggesting a safe harbor for 2017 whereby the fork is treated as a taxable event, but with a basis of 0, thereby deferring any tax payments until the asset is sold. This is similar to the way Japan treats hard forks. The Association of International Certified Professional Accountants (AICPA) suggests that taxpayers be allowed to make an election within 30 days of the fork, pay income taxes on the value in the election, and capital gains taxes would be applicable thereafter. Alternatively, if no election is made, all gains would be subject to income taxes when the asset is sold.

One of the most thorough analyses of the tax implications of chain splits to date is Landoni and Pieters (2019), in part because it went to press following the IRS Revenue Ruling 19–24 in October 2019. They identify several tax-related challenges that chain splits impose: parent/child designation (uncertainty regarding which is the continuation of the “original” chain), token access, and issues around fair market value. The challenges impact both determination of basis and timing of income. Landoni and Pieters assess the pros and cons of three frameworks to address chain splits: treasure trove, asset split, and calving. They conclude that calving, whereby the new coin is considered an offspring of the existing coin, assigned zero basis, and taxed upon sale, is the least problematic.

6 Intellectual Property Law

Regulators and legislators have long attempted to balance fairness, efficiency, social benefit, and commerciality as it relates to the ownership interests of individuals. A common legal theme throughout history has been to reward people for creating novel work by providing them a right to exclude others from deriving benefits without approval. As it relates to intangible property, intellectual property law emerged as a legal framework designed to prevent the dissemination of unauthorized reproduction that has the potential to erode a work’s commerciality (Zeilinger 2018).

6.1 Efficient Global Registration

In civil law, while copyright is typically granted upon the creation of work, other IP rights like patents or trademarks must be codified through a registration process. The process of registration for a right in intellectual property is typically complex, costly, requires a number of intermediaries, and only enables the IP holder to the rights in the registered country. Blockchains, and the technological benefits that come with using distributed ledger technology, has the potential to unlock a number of new benefits by making the registration process easier, faster, and cost-effective. In addition, the global nature of blockchain may also help overcome the issues surrounding different registration requirements across the world (Gürkaynak et al. 2018).

6.2 Provenance

Today, works in digital form are easy to download, copy, modify, and recirculate—potentially diminishing the commercial value of the intellectual property. Using an asset-aware blockchain like the open-source Ravencoin (RVN) project, intangible work could simply be ascribed to a blockchain's ledger, immutably recording its existence along with details of its provenance, value, and history. Once written to a blockchain, an authorized copy of the work could be identified as the original, and all other copies could be approved or prevented by the IP holder to circulate (Zeilinger 2018).

6.3 Preventing Counterfeit

For a work's commercial benefits to flow efficiently to a rights holder, the rights holder must be able to enforce his or her rights effectively (Gürkaynak et al. 2018). Since blockchains store all transactions across a shared database, distributed ledger technology has been identified as a promising solution in the prevention of counterfeit. The complex cryptography powering a blockchain empowers its immutability, irreversibility, and permanency. These benefits better allow individuals to track the ownership rights of their intangible property. In fact, there are several initiatives being developed today from firms like IBM that aim to better prevent counterfeiting (Pun et al. 2018). These blockchain-based solutions could play a critical role in enforcing IP rights in the future.

6.4 Challenges—First-Sale Doctrine

Implementing a blockchain-based system for IP law won't come without its challenges. It will require legislators and regulators to grant legal status to blockchains, develop acceptable legal standards for blockchain-based IP, and create efficient technological environments that best empower tracking and managing blockchain-based IP.

In addition, there are also a number of inherent conflicts that exist between existing laws and a new blockchain-based paradigm. For example, 17 USC §109(a), called the “first sale doctrine,” provides that the owner of a physical copy is entitled, without the authority of the copyright owner, to sell or otherwise dispose of the possession of that physical copy (17 USC §109(a)). Under this doctrine, a lawful owner of paintings, books, music albums, and memorabilia would be permitted to sell that physical copy without the express permission of the copyright owner (Fisher 2019).

However, digital copies are typically treated differently than physical copies under U.S. copyright's first-sale doctrine. In 2001, the U.S. Copyright Office published an opinion stating that a digital first sale right could not exist. In *Capitol Records LLC v. ReDigi Inc.*, the Second Circuit affirmed the U.S. Copyright Office's opinion and found that because it was impossible to transfer a digital file without making a copy, a transfer would be subject to a copyright owner's ongoing commercial reproduction right of digital work, as opposed to a copyright owner's distribution right of physical work (Fisher 2019). Under current U.S. law, the first-sale doctrine maintains that a work be tangible or physical in order to be relied upon.

Blockchain-based copyrights have the potential to blur these legal lines. NFTs, or non-fungible tokens, represent unique, authenticated, digitally scarce blockchain tokens. In addition to these digitally scarce tokens, smart contracts have the potential to create transferability controls on public blockchains like Ethereum. Under a blockchain paradigm, a digitally unique and scarce blockchain NFT has the potential to be transferred without resulting in a new copy, raising important new legal questions as to the legal definitions of “digital” and “physical” (Fisher 2019). This is one example of the types of legal issues that may arise as IP laws intersect with distributed ledger technology.

7 Concluding Remarks

At the most fundamental level, blockchains are a ledger of time-stamped transactions. Since there exist myriad types of transactions, or events, that might be recorded, there are numerous intersections with the legal and regulatory environment.

In this chapter, we scratch the surface of several of these intersections such as securities law, contract law, IP law, and tax law. However, as the types of information recorded on-chain expands, one could imagine an expansion of the law literature on this topic as well to areas like employment law, estate law, bankruptcy law, and perhaps even medical malpractice law. Ultimately, we expect the norms and customs that are being developed in these ecosystems to build the foundation of a new body of law, which Wright and De Filippi (2015) term *Lex Cryptographia*.

Although Bitcoin just passed 10 years old, we are still in the infancy of distributed ledgers. As the technology continues to advance, new use cases will emerge and adoption will continue to increase. While this unfolds, there will be substantial opportunities for additional scholarship at the intersection of blockchains and law.

Acknowledgements We thank Cameron Pfiffer and Clayton McDonald for research assistance.

References

- Barclay, Michael J., and Terrence Hendershott. 2003. Price discovery and trading after hours. *The Review of Financial Studies* 16: 1041–1073.
- Barclay, Michael J., and Terrence Hendershott. 2004. Liquidity externalities and adverse selection: Evidence from trading after hours. *The Journal of Finance* 59: 681–710.
- Bekaert, G. 1995. Market integration and investment barriers in emerging equity markets. *World Bank Economic Review* 9: 75–107.
- Finck, M. 2018. Blockchains and data protection in the European Union. *European Data Protection Law Review* 4: 17–35.
- Fisher, Katya. 2019. Once Upon a Time in NFT: Blockchain, Copyright, and the Right of First Sale Doctrine. *Cardozo Arts & Entertainment Law Journal* 37: 629.
- Foley, Sean, Jonathan R. Karlsen, and Tālis J. Putniņš. 2019. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies* 32: 1798–1853.

- Giancaspro, M. 2017. Is a 'smart contract' really a smart idea? Insights from a legal perspective. *Computer Law & Security Review* 33: 825–835.
- Goforth, Carol R. 2019. Securities treatment of Tokenized offerings under U.S. Law. *Pepperdine Law Review* 45: 405–470.
- Gürkaynak, Gönenç, İlay Yılmaz, Burak Yeşilaltay, and Berk Bengi. 2018. Intellectual property law and practice in the blockchain realm. *Computer Law and Security Review* 34: 847–862.
- Gurrea-Martínez, Aurelio and Nydia Remolina. 2018. The Law and finance of initial coin offerings. Ibero-American Institute for Law and Finance Working Paper No. 4/2018.
- Harvey, Campbell R. 2014. Bitcoin myths and facts. <https://dx.doi.org/10.2139/ssrn.2479670>.
- Henderson, M. Todd, and Max Raskin. 2018. A Regulatory classification of digital assets: Towards an operational howey test for cryptocurrencies, ICOs, and other digital assets. *Columbia Business Law Review*, Forthcoming.
- Jensen, M.C., and W.H. Meckling. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics* 3: 305–360.
- Landoni, Mattia, and Gina C. Pieters. 2019. Taxing Blockchain Forks. Available at SSRN: <https://ssrn.com/abstract=3475598>.
- Lerer, M. 2019. The taxation of cryptocurrency. Virtual transactions bring real-life tax implications. *The CPA Journal*. Retrieved from <https://www.cpajournal.com/2019/01/24/the-taxation-of-cryptocurrency/>.
- Lippiatt, Todd, and Michael Oved. 2018. The two token waterfall, Retrieved from: <https://www.tokenframework.io>.
- Maas, Thijs. 2019. Initial coin offerings: When are tokens securities in the EU and US? (February 13, 2019). Available at SSRN: <https://ssrn.com/abstract=3337514> or <https://dx.doi.org/10.2139/ssrn.3337514>.
- McKeon, Stephen. 2018. The security token thesis. Retrieved from <https://medium.com/hackernoon/the-security-token-thesis-4c5904761063>.
- Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G.M. Voelker, and S. Savage. 2013. October. A fistful of bitcoins: Characterizing payments among men with no names. In Proceedings of the 2013 conference on Internet measurement conference, 127–140.
- Metjahic, L. 2017. Deconstructing the DAO: The need for legal recognition and the application of securities laws to decentralized organizations. *Cardozo Law Review* 39: 1533–1568.
- Park, James J. 2018. When are tokens securities? Some questions from the perplexed. Lowell Milken Institute Policy Report (December 10); UCLA School of Law, Law-Econ Research Paper No. 18–13. Available at SSRN: <https://ssrn.com/abstract=3298965>.
- Pun, Hubert, Jayashankar M. Swaminathan, and Pengwen Hou. 2018. Blockchain adoption for combating deceptive counterfeits. Kenan Institute of Private Enterprise Research Paper No. 18–18. Available at SSRN: <https://ssrn.com/abstract=3223656>.

- Raskin, M. 2017. The law and legality of smart contracts. *Georgetown Law Technology Review* 1: 305–341.
- Rauchs, M., A. Glidden, B. Gordon, G.C. Pieters, M. Recanatini, F. Rostand, K. Vagneur, and B.Z. Zhang. 2018. Distributed ledger technology systems: A conceptual framework.
- Ritter, J.R. 1987. The costs of going public. *Journal of Financial Economics* 19: 269–281.
- Rogoff, K. 2016. *The Curse of Cash*. Princeton, NJ: Princeton University Press.
- Savelyev, A. 2017. Contract law 2.0: ‘Smart’ contracts as the beginning of the end of classic contract law. *Information & Communications Technology Law* 26 (2): 116–134.
- Solomon, Steven Davidoff. 2017. Dole Case Illustrates Problems in Shareholder System. *New York Times*.
- Soska, K., and N. Christin. 2015. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In 24th {USENIX} Security Symposium ({USENIX} Security 15, pp. 33–48.
- Szabo, Nick. 1996. Smart contracts: Building blocks for digital markets, U. AMSTERDAM, https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smartcontracts_2.html.
- Szabo, Nick. 1997. Formalizing and securing relationships on public networks. *First Monday*, 2. <https://firstmonday.org/ojs/index.php/fm/article/view/548/469>.
- Tapscott, Don, and Alex Tapscott. 2016. *Blockchain revolution: How the technology behind bitcoin is changing money, business, and the world*. Penguin.
- Van Slobbe, J. 2016. The drug trade on the deep web: A law enforcement perspective, In *Internet and Drug Markets*, EMCDDA Insights.
- Walch, Angela. 2019. In Code(rs) We Trust: Software developers as fiduciaries in public blockchains. regulating Blockchain. *Techno-Social and Legal Challenges*, edited by Philipp Hacker, Ioannis Lianos, Georgios Dimitropoulos & Stefan Eich, Oxford University Press, Forthcoming.
- Webb, N. 2018. A fork in the Blockchain: Income tax and the bitcoin/bitcoin cash hard fork. *North Carolina Journal of Law & Technology* 19: 283–311.
- Weekley, Roland. 2018. The Problematic Tax Treatment of Cryptocurrencies. *Florida State University Business Review* 17: 109.
- Werbach, K., and N. Cornell. 2017. Contracts ex machina. *Duke Law Journal* 67: 313–381.
- Wright, A., and De Filippi, P., 2015. Decentralized blockchain technology and the rise of lex cryptographia. Available at SSRN 2580664.
- Xu, D. 2018. Free money, but not tax-free: A proposal for the tax treatment of cryptocurrency hard forks. *Fordham Law Review* 87: 2693.
- Yermack, D. 2017. Corporate governance and blockchains. *Review of Finance* 21: 7–31.
- Zeilinger, M. 2018. Digital art as ‘monetised graphics’: Enforcing intellectual property on the blockchain. *Philosophy & Technology* 31: 15–41.

Zetzsche, D.A., R.P. Buckley, and D.W. Arner. 2018. The distributed liability of distributed ledgers: Legal risks of blockchain. *University of Illinois Law Review*, 1361–1406.



Blockchain Trading and Exchange

Hugo Benedetti, Stephen McKeon, and Cameron Pfiffer

1 Introduction

Blockchains are time-stamped ledgers of transactions or events, the contents of which are appended via a consensus mechanism, and secured through cryptography. Ledgers have been around for millennia, dating back to clay tablets in Mesopotamia. More recently, ledgers are often digitized using software such as Microsoft Excel, Quickbooks, and various other accounting software packages. Generally, ledgers are maintained by a single entity, which might be a bank for currency accounts, depository trusts for public securities, or a county recorder for real estate ownership. However, blockchains differ in this respect because the ledger is maintained by a software protocol running on a distributed network of hardware nodes, none of which have unilateral control.

Bitcoin, a public and permissionless blockchain, was envisioned as a peer-to-peer payment system, where anyone can transmit value to another party without the need for a third-party intermediary. However, to trade bitcoin for another asset, whether that be converting fiat currency into bitcoin or

H. Benedetti

ESE Business School, Universidad de Los Andes, Santiago, Chile

S. McKeon (✉) · C. Pfiffer

Lundquist College of Business, University of Oregon, Eugene, OR, USA

e-mail: smckeon@uoregon.edu

trading bitcoin for another cryptocurrency like ether, the transaction is often intermediated by a trading venue.

As we detail in this chapter, trading venues for cryptocurrency range from centralized exchanges like Coinbase and Bitstamp to more automated decentralized solutions like IDEX and Kyber Network. Additionally, traditional exchanges such as the Swiss Stock Exchange and NASDAQ Stockholm have listed products offering exposure to major cryptocurrencies like bitcoin. Finally, OTC desks process substantial volumes of cryptocurrency transactions.

Blockchains can also be used to trade, settle, and clear traditional assets like stocks and bonds. This may occur by integrating a blockchain into the back-office technology stack of existing venues, or creating a representational blockchain token for the asset that trades publicly across many different venues, or even peer-to-peer. A number of benefits have been proposed such as more efficient settlement, reduced trading costs, increased transparency around voting and governance, reduced regulatory compliance burdens, and interoperability across asset classes and jurisdictions.

However, a variety of challenges remain. Some of these are technical in nature, for example, cybersecurity attacks, network forks, and limitations to scaling. Others are economic and have been observed outside blockchain trading, but have yet to be solved in this new environment. Examples in this category include front-running, market segmentation, and undesired transparency. Finally, as the regulatory landscape struggles to keep pace with the rate of technological advancement, challenges along this dimension are present as well.

This chapter proceeds as follows. The first section reviews the landscape of trading venues and the second section examines other market infrastructure participants in the blockchain ecosystem. The third section synthesizes the literature around proposed benefits of blockchains for the trading process. The fourth section reviews challenges that must be overcome to realize the benefits. The final section concludes.

2 Landscape of Trading Venues and Other Market Participants

2.1 Overview of Blockchains

While a detailed description of the technology underpinning blockchains is outside the scope of this chapter, we begin with a broad description and some terminology.

The concept of time-stamped digital ledgers traces its roots to Haber and Storetta (1991) and one of the most famous applications is the Bitcoin network (Nakamoto 2008). Blockchain ledgers are distributed, meaning the contents are distributed across multiple parties. In the case of Bitcoin and many other blockchain networks, complete identical copies of the ledger are held by many different parties. Today, there are thousands of blockchain ledgers recording transactions involving not only cryptocurrency like Bitcoin, but also equities, debt instruments, real estate, commodities, and derivatives. Broadly, assets are often described as being “blockchain native” such as bitcoin or a lending contract where all underlying assets are cryptocurrency, or “traditional” where the asset is tied to something in the physical world like a corporation or real estate parcel. When an ownership claim on a traditional asset is recorded on a blockchain, the representation of this value is often referred to as a “token.”

New transactions are added through consensus. Effectively, this means that a majority of network validators agree on the legitimacy and order of the transactions that are posted. Transactions are typically added in batches, called blocks, and linked together through a cryptographic hash function, hence the name blockchain.

Blockchains can be public, meaning anyone is free to transact on the network and/or participate in consensus, or permissioned, where activity is restricted. Permissioned blockchains might designate specific entities to perform the consensus function (determining what transactions get added to the ledger) and only allow certain participants to transact on the blockchain, for example a group of banks. Public blockchains are often transparent, where all transactions are visible. Full transparency comes with both benefits and challenges, and technological advances such as Zero Knowledge Proofs offer enhanced privacy on public chains. Transactions on permissioned, or private, blockchains are typically only visible to admitted participants.

2.2 Trading Venues for Blockchain Assets

The most widely implemented type of trading venue is a centralized marketplace.¹ These marketplaces are managed by a single entity with full control of the design and operation of the trading platform, such as which assets are listed, user access and on-boarding, trading rules, order types, fee structure, and additional products (derivatives, margin trading, lending). Given the broad set of characteristics a marketplace can decide on, the market has evolved to accommodate widely differentiated sets of marketplaces along dimensions such as the number of tokens traded (with some trading over one thousand different tokens), order types (simple limit orders to sophisticated combinations of time to execute, slippage, hidden, and others), user identification and vetting (ranging from no registration up to banking grade KYC regulation), and token listing process and information disclosure (from no selection to full disclosure on selection process).

Benedetti and Nikbakht (2021) examines the cross-listing returns of tokens and finds higher cross-listing returns for tokens listing to crypto-marketplaces that provide access to new sets of investors, more sophisticated trading services, and a stricter regulatory environment. This finding points to the existence of a crypto-marketplace differentiation through “quality” and a market premium for such.

Ante (2019) uses a similar setting and analyses individual crypto marketplace listing returns and finds persistent positive abnormal returns for a reduced set of marketplaces, highlighting the effect of each marketplace choice of features on the market’s response to cross-listings.

One key feature on centralized exchanges relates to the centralized custody of all traded assets. This implies that in order to post a trade, a user must first transfer custody of funds to the exchange, by depositing cryptocurrencies, tokens, or fiat currencies in a compatible wallet/address/account. Once the funds are in custody of the exchange, a user account is created on an internal ledger that will reflect the funds balance. All trades performed within the exchange are recorded and settled on internal ledgers of the exchange (also referred to as off-chain records). There will only be a blockchain record (on-chain) when a user decides to withdraw funds to an external wallet/address/account. While some trading protocols allow centralized exchanges to perform on-chain recording and settlement of trades (e.g., Arwen), these solutions are still emerging as of this writing.

¹ Cryptocompare lists 220 centralized exchanges and 19 decentralized exchanges as of this writing: <https://www.cryptocompare.com/exchanges/#/overview>.

A variant on centralized exchanges are over-the-counter (OTC) desks. OTC is common in several traditional asset classes such as electronic swaps, where 95% of trades are OTC (Nagel 2016), and corporate bonds, where 81% of trades are OTC (McPartland 2016). In cryptocurrencies, OTC desks typically handle large block trades where market participants seek to avoid price impact. Major OTC desks for cryptocurrency include Amber AI, Genesis, Cumberland, and Circle. Additionally, brokers with smart order routers, such as Tagomi, promise best price execution by tapping into electronic price feeds from both centralized exchanges and OTC desks.

Decentralized exchange is a style of trading venue that is rising in prominence and is unique to blockchain trading. On a decentralized exchange, a software protocol (such as 0x or Kyber) performs the backbone of the trading process (order book building, matching engine and trading settlement), while the exchange determines asset listing and delisting, user onboarding and trading fees. The key difference is that on a decentralized exchange, trades are executed peer-to-peer and directly on-chain, as opposed to a centralized exchange, where the trades are executed off-chain and settled on the exchange's ledger.

In order to execute on-chain settlement, trades must be performed on assets that are either represented on the same blockchain protocol (Ethereum, WAVES, OMNI) or through a special software that connects different blockchain protocols and ensures the secure trading of assets, known as atomic swaps. While this feature eliminates the risk of hacks and thefts present in centralized exchanges, due to the fact that users trade directly peer-to-peer and the decentralized exchange does not hold any token in custody, it comes at the expense of publicly broadcasting the trading orders before they are posted on the blockchain. This transparency feature, inherent to most public blockchains, has led to front-running of orders as described in more detail in the challenges section of this chapter. A second burden of decentralized exchanges is the lack of direct fiat currency trading. As protocols only allow peer-to-peer trading of digital assets, fiat currencies need to be converted to a blockchain representation through tokenization. Tokenized fiat currency, often referred to as stablecoins, are discussed in the section covering other market infrastructure later in this chapter.

2.3 Securities Exchanges Integrating Blockchain Technology

The centralized and decentralized venues discussed above are largely trading assets that are not deemed to be securities, however, blockchains are intersecting with securities trading as well. There are several different approaches. For example, trading venues such as OpenFinance and tZero have acquired ATS licenses and focus on trading tokenized assets that are regulated as securities. Examples of traditional assets might include shares of equity, bonds, or ownership claims on real estate. A tokenized version of these assets simply creates a digital wrapper in the form of a blockchain token, so that ownership can be recorded and transferred on blockchains.

Another approach is where a traditional exchange integrates blockchain technology, which may mean a back-office solution for settlement and clearing, or a traditional security format that offers exposure to cryptocurrency, or full support for tokenized trading.

MERJ, a stock exchange located in Seychelles, created a framework for digital assets which allows it to list, trade, clear, settle, and register crypto securities within a regulated infrastructure. In August 2019, MERJ listed the first tokenized security on a stock exchange globally. To date, MERJ is the first regulated market to facilitate primary and secondary markets for digitized securities.

GBX, parent company of Gibraltar Stock Exchange, created the Gibraltar Blockchain Exchange in 2017, aimed at developing the technological infrastructure and technical expertise required to implement blockchain technology in the traditional security trading process. As a first stage, they introduced Digital Asset Exchange (DAX), a crypto currency trading platform compliant to the regulatory framework of the Government of Gibraltar and GBXGrid, a token launch and placement platform. The second stage is to develop a token issuance framework that would allow companies to issue digital assets and simultaneously list on the traditional Gibraltar Stock Exchange.

The Swiss Exchange (SIX), created SIX Digital Exchange with the aim to be “the first market infrastructure in the world to offer a fully integrated end-to-end trading, settlement and custody service for digital assets. The service will provide a safe environment for issuing and trading digital assets, and enable the tokenization of existing securities and non-bankable assets to make previously untradeable assets tradeable.” Additionally, Amun has created several exchange traded products on SIX offering investors exposure to a variety of cryptocurrencies. The regulatory environment in Switzerland

has proven to be accommodating to innovation around cryptocurrencies. In contrast, the Securities Exchange Commission in the U.S. has rejected numerous requests for Exchange Traded Funds containing cryptocurrency.

In July 2018, the Santiago Stock Exchange implemented a permissioned blockchain to allow tracking and settling short sales. The blockchain infrastructure reduced the average back-office processing time from 5 days to 90 seconds and increased short sale volume ten-fold in less than one year. In a joint venture with the Chilean Central Securities Depository (DCV) and a local telecom (GTD), the group is developing a permissioned blockchain to serve as the digital backbone for blockchain developments in the Latin-American financial sector. Lastly, the Chilean Central Bank and the DCV are preparing the issuance of the world's first blockchain native Central Bank Bond. The issuance will not be a tokenization of a bond, but a direct blockchain issuance, with no "physical" counterpart.

Boerse Stuttgart launched a smartphone application (BISON) to enable cryptocurrency trading. A financial services provider, subsidiary of Boerse Stuttgart (EUWAX AG) operates as the trading partner and is the counterparty to all cryptocurrency trading operations.

Finally, the Australian Securities Exchange (ASX) announced the replacement of its current Clearing House Electronic Subregister System (CHES) with a new system based on blockchain technology (Chohan 2017). According to ASX, this implementation would create advantages in efficiency, delivery, privacy, interoperability, accessibility, and reliability.

3 Beyond Trading Venues: Other Market Infrastructure in Blockchain Trading

3.1 Fiat-Equivalent Assets On-Chain (i.e., Stablecoins)

In order to realize the full potential of blockchain based trading, where intermediation is minimized to nothing more than an automated smart contract, both assets must be on-chain. As discussed above in the section on Decentralized Exchanges, this is currently possible when trading two cryptocurrencies that reside on the same blockchain, such as two tokens on Ethereum. There are also some early examples of enabling trading across two different chains, known as atomic swaps, or atomic cross-chain trading. However, the trades in the largest traditional financial markets today typically exchange an investment asset for fiat currency. Therefore, in order to facilitate trades entirely with blockchains requires an asset on-chain that represents fiat currencies like

US Dollars. These assets are known as “stablecoins.” Given the importance of stablecoins for blockchain trading, we briefly review the landscape around this asset type.

Stablecoins are a class of cryptocurrencies that exhibit very low price volatility relative to a predefined benchmark. The benchmark is typically a fiat currency like USD, but could also be a price index, such as the consumer price index (CPI). There are two main categories of stablecoin models that have been widely deployed: fiat-backed and crypto-backed, discussed below. We note here that necessary conditions for widespread adoption of a stablecoin are that it is: (i) stable against the designated benchmark in virtually all states of the world, (ii) liquid and/or redeemable in virtually all states of the world, and (iii) easily scalable in both directions (expanding and contracting supply) without violating (i) and (ii).

Real asset-backed: Tether is currently the largest asset-backed stablecoin. As of this writing, over \$4 billion USD Tether are outstanding making it by far the most well established stablecoin. Tether asserts that they are backed 1-to-1 with assets held in reserves, however, the reserves include assets other than fiat currency, as well as receivables from loans made by Tether to third parties, which may include affiliated entities. The takeaway is that holders of Tether, and any other real asset-backed stablecoin, are exposed to counterparty credit risk, but are not rewarded for bearing this risk. Griffin and Shams (2019) present evidence suggesting that Tether is printed unbacked and pushed to the market to inflate the price of Bitcoin, particularly by a single trader. Tether disputes the findings, as have others in the industry.² To provide an alternative to Tether, many other fiat-backed stablecoins have emerged recently such as USDC, Paxos, Fidelity (USC), JPM coin, and Libra.

An important, but unsettled, issue around stablecoins is the degree to which they must adhere to Know Your Customer (KYC) and Anti-Money Laundering (AML) laws. First, lack of KYC makes it hard to establish banking relationships, which has been an ongoing challenge for Tether. Second, lack of KYC subjects the issuer to regulatory risk, which is related to the first point, but also broader. Holders can't trust the stablecoin to be stable if a regulatory body can freeze the underlying reserves. The most important and most criticized feature to understand about fiat-backed stablecoins is that they are not decentralized by construction and require a trusted third party, which is often cited as a critical point of failure.

Crypto-backed stablecoins are those in which there is no link to the traditional financial system, but the token has backing in the form of other

² See <https://tether.to/tether-response-to-flawed-paper-by-griffin-and-shams/>.

blockchain assets. They are less susceptible to counterparty risk and asset seizure, because there is typically no asset cache that resides in the traditional financial system as in real asset-backed stablecoins, but other risks are present. Some observers doubt these projects will succeed.³

Most crypto-backed stablecoins employ a multi-token set-up. The first is the stablecoin and the second (or third) is a risky asset that provides the collateral backing for the stable asset. Prices are managed through open market operations, which is to say as the price drops below the benchmark, then the market participants are incentivized to commence buying until it returns to \$1.00, and vice versa. MakerDAO is currently the largest project of this type and accepts several cryptoassets as collateral.

There exist other stablecoin models based on seigniorage, and the degree to which central banks will directly issue digital currency remains an open question. As these and other models evolve, they are likely to serve as trading pairs with cryptocurrency and other tokenized assets.

3.2 Base Layer Protocols

Base layers refer to the underlying blockchains upon which assets are issued. The base layer is typically associated with a native token, for example ether on Ethereum, which is used to pay transaction costs. Other assets can issue tokens on top of the base layer, with a separate ledger of ownership—the base layer is simply used as a service to maintain the ledger. When transfers occur, transaction fees are typically paid in the base layer's native asset. These blockchains are often referred to as “smart contract platforms” because the economic activity on chain is regulated by smart contracts.

As of this writing, Ethereum is by far the dominant base layer smart contract platform. Coinmarketcap.com reports that 44 of the top 50 tokens by market cap use Ethereum as the base layer. Other base layer blockchains that are targeting tokenization of financial securities and other assets include Algorand, Stellar, Tezos, Ravencoin, Findora, Ownera, and Polymesh.

3.3 Clearinghouses

Post-trade clearing and settlement for many traditional financial securities is currently managed by large central securities depositories such as The Depository Trust & Clearing Corporation (DTCC), Euroclear and Clearstream. As

³ See, for example, <https://prestonbyrne.com/2017/12/10/stablecoins-are-doomed-to-fail/>.

blockchain-based trading grows these firms are directly impacted. In 2016, DTCC posted a whitepaper outlining their approach to distributed ledgers and more recently selected a consortium of IBM, Axoni, and R3 to implement a pilot for credit derivatives with several large financial institutions.⁴

3.4 Tokenization Platforms

A variety of firms facilitate primary issuance of blockchain assets. For example, Securitize has brought several tokenized securities to market including Blockchain Capital, a tokenized venture capital fund, Aspen Digital, a tokenized claim on real estate, and Lottery.com, a revenue sharing token. Nivaura is working with banks such as Banco Santander to tokenize traditional bonds. Other security token platforms include Harbor, Tokeny, Polymath, KoreConX, Swarm, and Neufund. To date, most of the issuances from these firms have occurred on the public Ethereum blockchain. Additionally, a number of these firms have acquired transfer agent licenses and registered as broker-dealers.

Concurrently, several firms have created platforms to tokenize assets on private or permissioned blockchains. Examples include Symbiont's Assembly blockchain and Figure's Provenance blockchain.

Finally, we are also seeing the emergence of tokens representing real world commodities. In September 2019, Paxos announced an Ethereum token backed by, and redeemable for, physical gold. The underlying assets are stored in professional vault facilities in London and the issuance was approved by New York State Department of Financial Services (NYDFS).

3.5 Derivatives

There are a variety of derivative products related to blockchains, ranging from cryptocurrency derivatives trading on traditional exchanges, to natively digital derivatives trading on both regulated and unregulated crypto exchanges.

The first major exchange to offer regulated futures trading on bitcoin was Chicago Mercantile Exchange, in December 2017. CME Group recently announced that it intends to add bitcoin options in 2020. Other notable new entrants addressing regulated cryptocurrency derivatives in the US are Bakkt, a project affiliated with ICE, and LedgerX. In Europe, Quedex

⁴ See <https://www.dtcc.com/news/2017/january/09/dtcc-selects-ibm-axoni-and-r3-to-develop-dtccs-distributed-ledger-solution>.

recently received approval from BaFin, the German financial regulator, to offer derivatives on cryptocurrency.

However, prior to the emergence of regulated derivatives, a market developed through less regulated channels. Alexander et al. (2019), study the role of the largest unregulated derivatives trading platform (BitMEX) on the price discovery and informational efficiency of Bitcoin. BitMEX is an unregulated crypto exchange operating from Hong Kong and incorporated in the Republic of Seychelles, offering highly levered derivative contracts on Bitcoin, with an average daily trading volume of over 2 billion dollars. The authors find that BitMEX derivatives lead prices in the major spot crypto exchanges. Moreover, BitMEX is informationally more efficient than the major spot crypto exchanges. Patti (2018) discusses some of the risks for institutional investors investing in Bitcoin derivatives offered in US regulated derivatives markets. Investors participating in these regulated instruments, receive exposure to the underlying through designated market makers and swap execution facilities, reducing the risks associated with trading bitcoin in unregulated exchanges. However, the reliability and accuracy of the current mechanisms for determining the reference price have not been fully documented.

The primary use-case for cryptocurrency derivatives has historically been speculation, but hedging is becoming an important motivation as the ecosystem develops. Lending of cryptoassets has grown substantially over the past year with roughly \$500M currently locked in automated lending applications that use smart contracts to mediate the positions. This segment of the ecosystem is known as “Decentralized Finance” or DeFi. Notable participants in the DeFi market for lending are MakerDAO, Compound, Synthetix, dYdX, Nuo, and Dharma. Additionally, centralized lenders include BlockFi, and Genesis Trading. In most cases, the lending is done on a secured basis, where one type of cryptocurrency is pledged as collateral for the loan. If the value of the collateral drops below a predefined threshold, it risks being repossessed and sold by the lender. Derivatives are increasingly being used to offset this risk.

4 Potential Benefits of Blockchain Trading

As described above, the landscape of blockchain trading venues and associated market infrastructure is developing at a rapid pace. This rising tide of new market infrastructure promises numerous benefits. In this section, we focus

on a subset of proposed benefits that have attracted the most attention, specifically, settlement, trading costs, voting/governance, regulatory compliance, and interoperability.

4.1 Settlement and Clearing

A commonly cited benefit of blockchains in financial markets is their application in the post-trade processing system. After a trade on a traditional securities exchange, many parties must interact to verify that any security traded is one that has previously been issued, that all parties are authorized to trade, and that the trade actually occurred. Broadly, the extant post-trade processing industry involves three major functions (Benos et al. 2017):

1. Order management. This is the process by which transaction information is recorded and validated.
2. Clearing. Clearing is the process of netting out transactions between parties, calculating margins, and novation.
3. Settlement. At the end of the post-trade process, a Central Securities Depository (CSD) facilitates the exchange of securities and cash between parties.

Some or all of these functions could be replaced by a distributed ledger, which would simultaneously standardize industry data practices and reduce the need for a single trusted counterparty like the CSD. A deeper examination how settlement functions could be implemented in a distributed ledger can be found in Mills et al. (2016). A distributed ledger as applied to the post-trade settlement system could increase processing speeds, reduce risk, and decrease costs.

The most utopian vision might be that of Micheler and von der Heyde (2016), where the entire system of trading, payment, clearing, and settlement are all unified into a single monolithic distributed ledger (or many linked ledgers). All cash and financial assets would be recorded on the blockchain in the system they envision. Caytas (2016, p. 10) suggests that such a system would trivialize much of the current post-trade processing industry: “Digitization of securities custody, delivery, and payment reduces the clearing and settlement processes largely to a software issue.”

The advantages of an entirely distributed post-trade system are numerous. The unified nature of trading would allow for settlement to occur at a rapid pace. Regulators would have high visibility into market activity. The need for custodians and various other intermediaries would be obviated, along

with the custodian risk that accompanies it. Further, counterparty risk would be drastically reduced by the use of smart contracts that rigidly enforce fulfillment.

Pinna and Ruttenberg (2016) present a more balanced review of the degree of integration a distributed ledger might have with different layers of the post-trade process. They suggest that the notary function (ensuring that invalid securities are not created) must still be performed by a centralized third party, but that the validation of transactions and the transfer of ownership could be built onto a distributed ledger. The settlement layer could be made more robust by allowing the ledger to track cash positions as well (i.e. stablecoins), and thus a security and the cash used to purchase it could be exchanged immediately.

Settlement times and transaction costs could be substantially reduced. Chiu and Koppl (2019) suggest that average settlement times in the US corporate debt market could be reduced from two days to 148 minutes on average with the use of a permissionless distributed ledger. Additionally, average transaction costs could be reduced by 1–4 basis points.

A challenge within settlement is that transactions on a blockchain are not final in the same sense as in traditional markets. Pinna and Ruttenberg (2016) and Benos et al. (2017) highlight that finality on a blockchain is fundamentally probabilistic: a transaction becomes more “final” as more blocks are added to the chain and it becomes increasingly difficult for an adversarial participant to alter the ledger, thereby invalidating the transaction. Probabilistic finality poses a variety of legal issues regarding settlement that are outside the scope of this paper, but interested readers can refer to Benos et al. (2017) for more detail.

The custody layer could be removed entirely. Market participants can directly hold their own assets on a distributed ledger instead of a custodian holding the assets on an investor’s behalf, and smart contracts would allow for the automatic distribution of dividends and interest payments directly to the investor. Such a system may result in significant cost savings to investors, as revenue earned by custodians neared \$39 billion in 2014 (Oliver Wyman/Swift 2014). The custody layer may be the most likely to be implemented first, since institutions can agree to build distributed custody systems to eliminate the fees they pay to custodians. That said, some institutions require custodians for legal or regulatory reasons so evolution along these lines will be required for widespread institutional adoption of self-custody.

As for the clearing layer, the outcome of distributed ledger applications is less obvious and are dependent upon the degree of integration between the trade and settlement system. If trades and settlement occur on the same

ledger, then much of the traditional work involved in securities clearing such as order reconciliation, netting, and novation become unnecessary. As an example, trading on a distributed ledger allows near instantaneous settlement in comparison to traditional $T + 1$ or $T + 2$ settlement cycles. Traders would be required to have cash and securities on hand before making a trade, and counterparty risk would be eliminated since securities and cash would be exchanged at the same moment with no risk of default. Instantaneous settlement essentially eliminates counterparty risk and the usefulness of novation, but it introduces the requirement that firms have cash on hand at the time of a trade (Khapko and Zoican 2019). Netting would not be required since all trades are cleared immediately.

Clearing functions could be replicated on a distributed ledger even if settlement is not instant, though the lack of instantaneous settlement introduces counterparty risk. Counterparty risk could be reduced on a distributed ledger the same way it is currently managed: margins could be required to compensate a counterparty in the event of a default, and netting can reduce outstanding positions. Smart contracts could net positions between counterparties, and margin calls can be handled automatically if the distributed ledger is able to access traders' collateral.

The benefits to capital market operational efficiency go hand-in-hand with skepticism. Mainelli and Milne (2016) frame the discussion of the distributed ledger settlement process around its feasibility. They conduct interviews and focus groups with practitioners in both the post-trade processing and distributed ledger industries. Mainelli and Milne note that the blockchain has the potential to reduce cost and risk, but that there are significant expected up-front expenditures to transition the current post-trade processing industry to a distributed system.

4.2 Trading Costs

Trading on a distributed ledger has the potential to reduce transaction costs faced in traditional markets, while introducing new costs that are not well studied. Many blockchain-based (both centralized and decentralized) exchanges feature all the explicit and implicit costs that appear in traditional trading markets, and include maker/taker fees, withdrawal/deposit fees, adverse selection costs, and inventory costs. These are well studied outside of the blockchain literature, however, trading systems built onto a blockchain, or those that otherwise interact with blockchains, introduce new considerations for academics, regulators, and practitioners. New explicit cost categories

include fees paid to miners, and new implicit costs include transaction processing time and the cost of pseudonymity.

Validators on blockchain networks are typically compensated through two mechanisms. One is a block reward, where small amount of new cryptocurrency is minted and paid to miners that post blocks to the ledger. These fees are indirect in the sense that the cost is borne by all cryptocurrency owners through inflation rather than the traders. Additionally, most cryptocurrency blockchains allow (or require) a discretionary fee to be included with transactions. Users can include a higher fee to incentivize miners to include a transaction more quickly, or they can use the default baseline transaction fee which provides no guarantees as to the processing time for any transaction. In some cases, they can choose to post no fee at all.

These fees are largely determined by the platform upon which a trading mechanism is constructed, and evolve over time and technical design. Fees are not typically based on value of the transaction, but the complexity of the transaction (for example, multiple signatures or addresses) can impact fees. In some cases, fees may be required to reduce spam transactions to increase the security of the ledger, while in permissioned systems fees may not be needed at all.

For example, Ripple's XRP token, which is designed primarily for institutions, requires a fixed fee for any asset transfer of 0.00001 XRP, which as of September 17th, 2019 is worth approximately \$0.000003. This fee can increase or decrease in response to the aggregate number of transactions to compensate validators for larger volumes.

Fees can impact the time it takes for a transaction to be posted to the ledger. Less risk-averse users may use the lowest possible fee without caring when or if their transaction is processed, while firms who require immediate settlement may pay in accordance with their risk tolerance. Chiu and Koepl (2019) model waiting times and transaction processing speeds jointly, suggesting that transaction costs in a proof-of-work blockchain should increase with information arrival rates and the average time it takes to add a new block to the chain ("block time").

In a similar vein, Easley, O'Hara, and Basu (2019) model the interactions between waiting times and transaction fees on the Bitcoin blockchain. Easley et al. find that the percentage of transactions paying zero fees decreases with waiting times, but that the level of average transaction costs are not significantly related to waiting times. As they note, this is counterintuitive—one possible reason is that users who submit transactions with zero fees are typically a highly specialized group who have no time preference over when their transaction is processed.

Malinova and Park (2017) address the costs of trading on a decentralized blockchain network when anonymity is not guaranteed. On most existing permissionless blockchains, balances and transaction history for each address are fully transparent, in contrast to traditional order-driven markets where traders do not have any information about the counterparty. The inclusion of this new information has the ability to impact the behavior of traders, which can be considered in the same vein as other implicit costs like asymmetric information. Malinova and Park (2017) find that aggregate welfare is improved when traders are allowed to mask their total holdings by splitting their assets across many addresses, which cannot be linked to their ultimate owner. In their model the main cost is the potential that a liquidity demander may be front-run, but future work may examine in more detail what traders are able to infer using publicly available information on addresses. A variety of analytics providers such as IntoTheBlock and Coin Metrics offer detailed data for empirical analysis.

4.3 Proxy Voting/Governance

Shareholder voting presents challenges to firms from a logistical and corporate governance standpoint. Shareholders of corporations typically have the right to vote on key corporate decisions, which may include representation in the board of directors, management compensation, capital structure decisions, and/or mergers and acquisitions. The voting process is usually held during shareholders' meetings, where each shareholder holding voting stock is allowed to attend and cast their vote. For firms with large shareholder bases, several logistical problems arise from this process. First, how to properly account and track the identity of all shareholders; second, how to provide information regarding the voting to all shareholders in a timely manner; third, how to securely allow non-presential voting or presential through representatives ("proxy-voting"); and lastly, how to provide an end-to-end auditable anonymous voting process.

These problems are accentuated by the fact that most retail investors are not direct shareholders, but indirectly own securities through investment vehicles such as mutual funds, and pension funds. In this case and depending on the applicable regulation, retail investors are given the choice of participating directly in the shareholder's meeting and casting votes; or indirectly through a representative, usually the investment vehicle administrator.

These challenges might be improved by tokenizing or issuing shares using blockchain infrastructure. In this scenario, total share issuance could be

visible and trackable in real time. Depending on the identification requirements, each token holder address could also be directly or indirectly linked to an individual investor. Information issued by the company could be distributed simultaneously and publicly to all token holders through the blockchain network and votes could be cast securely, anonymously and in an auditable manner using current encryption technology.

Kahan and Rock (2007) provide a background of the key problems with proxy voting, highlighting the relevance of accurate share balance and vote counting, given the fact that some controversial corporate decisions subject to shareholder voting have been passed by narrow majority margins, such as the Compaq and HP merger with 51.4% of shares, AXA/MONY merger with 53.8% of shares, and the Transkaryotic merger with just 52% of shares.

More recently, Fox (2017), highlights the need of electronic databases, such as distributed ledgers or blockchains, to improve not only the voting count, but the official share count itself, citing the textbook example of proxy voting failures: Dole Food Co's take-private deal, where 49.2 million shares cast votes, despite only 36.8 million shares of the voting stock outstanding in the company's registry.

McCorry et al. (2017) provide an implementation of a decentralized and self-tallying Internet voting protocol, using an Ethereum smart contract. Their implementation provides full voter privacy, allows voting outcomes to be displayed publicly, and does not require the intervention of a central authority.

In contrast, Bramhall (2019) analyzes the implications of tokenizing equity securities and argues that proxy voting solutions offered by blockchain technology can be achieved by currently available systems, and that the challenges arising from using blockchain technology far outweigh the potential benefits.

4.4 Regulatory Compliance

Given the evolution of the financial market, and heightened by the financial turmoil of the last decade, regulatory agencies around the world have brought forward a large number of policies and regulations, such as Dodd-Frank, European Market Infrastructure Regulation (EMIR), Markets in Financial Instruments Directive (MiFID) I and II, Markets and Financial Instruments Regulation (MiFIR) and Securities Financing Transactions Regulation (SFTR), among others. In general, these regulations aim to improve financial market infrastructure by increasing market transparency (pre and post trade), reporting, disclosures, market monitoring, and record keeping.

As presented by Paech (2016), blockchain financial networks hold three characteristics (distributed ledgers, immutability of records, and the possibility to create smart contracts) that could disrupt traditional market practices and governance. The distributed ledger infrastructure potentially provides each participant with the complete record (or evidence of the existence and integrity of the complete record) of transactions, therefore “distributing” the record keeping process from current central parties to the full network. The immutability of records reduces the risk arising from unauthorized amendments performed by record keeping intermediaries. Smart contracts allow the automatic execution of actions according to a set of predetermined conditions, without human intervention or discretion.

Peters and Vishna (2018) highlight the benefits of permissioned blockchains. In particular, the potential to validate the identity of each participant and determine their entrance and removal from the network. Moreover, their closed environment allows a more purpose-built configuration, tailoring the blockchain capabilities to complement and link with current financial infrastructure. Peter and Vishna (2018) also describe the potential of smart contracts (self-enforcing contracts executed on-chain) to improve and automate regulatory compliance of transparency, transaction execution, and reporting.

Ducuing (2019) discusses the notion of “code is law” and how blockchain infrastructure allows the decentralization of certain activities, in particular by the use of transparent contracts coded on the blockchain that can be enforced without the need to resort to judicial authorities. The author discusses that while the open code/open source ethos of blockchain infrastructure provides ground for “trust” and confidence in the enforceability of smart contracts, it’s not able to generate a “meta-trust,” that is, to trust in the completeness and correctness of the blockchain infrastructure itself. Therefore, the potential virtues of smart contracts need to be complemented with improvements on the traditional regulatory framework to validate and enforce their legal execution.

Wright and De Filippi (2015) coined the term *Lex Cryptographia* to refer to rules administered through self-executing smart contracts and decentralized autonomous organizations, highlighting the benefits in terms of security, speed and transparency of removing intermediaries and decentralizing not only contracts, but markets, laws, and judiciaries. This study and several others that touch on the legal and regulatory environment are described in more detail in the Law and Blockchain chapter of this handbook.

4.5 Interoperability

Many of the benefits ascribed to blockchain trading revolve around interoperability, or in other words, the ability of different types of software and hardware to exchange and interpret information. Currently, the software utilized by global financial markets is not fully interoperable. For example, different brokers, custodians, and traders can access US equity markets from myriad software platforms, however, other markets, such as equities in Asia or commodities in the US, run on different systems. Blockchains are not required for interoperability—DTCC and Euroclear are already working on integration of collateral processing through a joint venture called Global Collateral Ltd. However, blockchains can create a common set of open standards upon which any developer can build interfaces. Ethereum's ERC-20 standard is one of the key innovations that led to the ICO boom.

Widely adopted standards at the protocol level facilitate interoperability, but interoperability can also be achieved at higher layers of the technology stack. For example, Cosmos is working on facilitating transfers and communication across blockchains. As these solutions develop, they are likely to impact segmented markets both across jurisdictions and across asset classes.

5 Challenges to Blockchain Trading

The aforementioned benefits have yet to be realized in most cases and a variety of challenges to implementation exist. In this section we review economic challenges such as front-running, market segmentation, regulation, and undesired transparency as well as technical challenges such as cybersecurity attacks, network forks, and limitations to scaling.

5.1 Front-Running

Front-running refers to an adversarial trader using advance information about a trade to profit at the expense of the trade initiator. For example, if a mutual fund needs to purchase 100,000 shares of a stock, it may distribute the purchase orders across multiple exchanges. A high-frequency trader can observe the trade in one market before the orders arrive to the remaining four markets, and place orders on those markets to purchase securities ahead of the mutual fund, only to resell them at a higher price when the mutual fund's order arrives.

Decentralized exchanges are not immune to this, and are especially vulnerable to front-running. Trades are not added to the blockchain instantaneously. After a trade is matched, the transaction is added to the pool of pending transactions. It is possible for an attacker to observe a trade in the pending transactions, initiate a transaction prior to the original transaction, and then pay a higher fee to get the front-running order processed first.

Eskandari et al. (2019) provide a taxonomy of front-running attacks that are currently possible on Ethereum. The first type is a displacement attack, where an adversary observes another person's intention from an Ethereum transaction and issues their own transaction before the first, displacing the original transaction from completing in the expected manner. The second type is an insertion attack. In an insertion attack, an adversary might observe a bid for an asset, and then publish their own transaction which buys the security at the existing price, and then immediately fills the original order. The profit of the spread goes directly to the attack. The third type is a suppression attack, which is simply intended to delay a good-faith transaction.

Daian et al. (2019) study insertion attacks on the Ethereum blockchain in great detail. They demonstrate that algorithmic traders capture between \$10,000 and \$100,000 per day using this front-running strategy. There is even a competitive market for arbitrageurs who engage in priority gas auctions (PGA), bidding their transactions fees ever higher to be the first to perform an insertion attack.

One crucial difference between traditional exchanges and a distributed ledger trading system is that, in a traditional exchange, the rules are fixed. Time-price priority cannot be changed and is largely binding: an order that arrives first has priority when a matching trade arrives. All traders' behavior is formulated on the basis that the rules by which the exchange operates are fixed. In a blockchain, this is not necessarily the case. Traders may circumvent the rules of time-price priority by ordering transactions as they see fit (Aune et al. 2017), submitting false transactions to delay the processing of legitimate transactions, or to reverse transactions within a narrow window of when they occurred by spending significant computational resources.

Lastly, blockchains often have greater latency compared to traditional markets. Latency is defined as the time between when an order is generated by the trader and when it arrives at the exchange. Traditional markets have extremely low latencies: in the microsecond or millisecond range (O'Hara 2015). In a distributed ledger setting, it is less clear what latency means. Latency may be the time when an order enters the pool of pending transactions, but all miners observe such transactions at different times and speeds.

It may be when a transaction is included in a block, which can take several seconds or minutes depending on the type of blockchain used. Regardless, as of this writing, latencies are universally higher in blockchain systems. Potential solutions might be implementing frequent batch auctions (Budish et al. 2015) or by using a permissioned blockchain with more rapid consensus.

5.2 Transaction Finality, no Recourse

Transactions on a blockchain are often characterized as immutable, in that they typically cannot be reversed and there is no third party to correct unlawful or erroneous behavior. If your private key is compromised and your funds are stolen, you cannot get your funds back unless the attacker returns them. There is no ex post settling up function to reverse a mistaken transfer. Transaction finality has had severe effects in the case of hacks on centralized exchanges, and in cases of bugs or vulnerabilities with the software used to interact with the blockchain.

Hacks and theft are attractive crimes, since funds cannot be forcibly returned and the inability to determine the ownership identity of a given address makes traditional legal solutions difficult to implement. Moore et al. (2018) study 80 cryptocurrency exchanges founded between 2010 and 2015. In their sample, 25 exchanges suffered a security breach, and 15 of those subsequently closed. Moore et al. (2018) find that suffering a breach has significant predictive power for whether an exchange will close in a given quarter.

The prevalence of the attacks on centralized exchanges is largely due to the fact that they custody what are effectively bearer assets on behalf of customers. These exchanges manage wallets on behalf of their clients, and thus all wallets at an exchange share common security vulnerabilities. Fully decentralized exchanges do not suffer from the risk of being hacked in the same way, because the decentralized exchange does not custody users' private keys and does not have access to client assets. Some emerging solutions such as Arwen allow for the speed of a centralized exchange while allowing users to hold their own assets, which reduces the risk to users if the exchange is compromised.

Transaction finality effectively creates largest bug bounties for those that discover and exploit flaws in the code and there are several examples of smart contract hacks. In June 2016, a hacker was able to drain roughly \$70 million in ether from The DAO smart contract. In 2017, a bug was discovered in Parity, a popular tool for mining and wallet management on the Ethereum blockchain, where nearly \$30 million worth of Ethereum was stolen from

three smart contracts. We hasten to note that the Ethereum protocol itself was not compromised in these cases, the vulnerabilities were due to faulty smart contracts running on top of Ethereum. In the case of The DAO, transaction finality was violated through execution of a hard fork, which created a new version of the transaction history that did not include the hacked transactions.

Bugs are unfortunately a common feature of software, but in a smart contract, which has dominion over assets, they can come at great cost. Smart contract auditing is a rapidly developing service category, examples of firms in this vertical include Zeppelin and Runtime Verification.

5.3 Attacks: DDoS, Goldfinger, 51%

Blockchain-based systems (particularly permissionless systems) are typically open by design, which means they are more subject to attack from a technical perspective. Blockchains are resilient to attacks to a degree by making it economically infeasible to act in bad faith. This resiliency does have limits, however, and there are cases where a blockchain can be compromised. Common attack vectors include 51% attacks, distributed denial of service attacks, and Goldfinger attacks.

In a 51% attack an adversary obtains control over a majority of the mining resources. Adversaries can then ignore legitimate transactions or engage in double spending in order to extract wealth from the blockchain. 51% attacks are extremely costly on chains with large amounts of hashpower, and are thus unlikely to be able to net an economic profit for an attacker (Kroll et al. 2013).

Economic profit may not be the primary goal, however. Blockchains may also suffer from a Goldfinger attack (named after the eponymous Bond villain) where an adversary compromises a blockchain for reasons other than profit. One example might be a foreign government which targets a domestic government's foreign exchange trading platform in order to destabilize exchange rates. Bonneau (2019) studies the risk factors involved in a Goldfinger attack on a blockchain network. Bonneau (2019) suggests that proof-of-stake and proof-of-work protocols are both susceptible to at least one form of Goldfinger attack by an entity with enough resources.

A distributed denial of service (DDoS) attack is an attempt to interrupt or slow the normal functioning of a system by sending a large amount of messages. In blockchains, this can occur in two forms.

First, the mining pool may suffer from a DDoS attack, where an attacker injects a high volume of faulty or expensive transactions designed to force the

miners to expend time and energy to locate valid transactions. One notable example of this is the DDoS attack on the Ethereum chain in September 2016, where an attacker was able to dramatically slow the processing of new transactions (Buterin 2016). Johnson et al. (2014) provide an analysis of the strategic application of a DDoS attack on a mining pool. A DDoS attack conducted on the mining pool is also an attack on every system that depends on the chain, and can severely impair the function of a blockchain.

Second, an endpoint of a chain may also suffer from a DDoS attack in the same way that any service that interacts with the Internet might be attacked. This includes any centralized or decentralized exchange, distributed application, or any other service that accepts some form of digital message from a user base. DDoS attacks on endpoints are not a unique issue to blockchains, and thus will not be considered in greater depth here.

Blockchains include many of the risks that traditional technical systems are exposed to, with several new considerations. Policy makers, practitioners, and academics should be wary of the fact that any implementation of a distributed ledger trading system requires a careful analysis of actions that can be taken by adversaries, whether motivated by profit or another purpose.

A comprehensive view of the operational risks in a blockchain trading system can be found in Walch (2015), who ultimately suggests that permissionless blockchains may not possess the appropriate level of reliability to conduct any large-scale financial activity. Walch's analysis is limited to Bitcoin's blockchain, though other researchers have arrived at a similar conclusion that trading systems will likely be a permissioned system (Mills et al. 2016), which is less likely to suffer from the faults highlighted in Walch (2015).

5.4 Segmented Markets—Price Dispersion

While cryptocurrencies and blockchain assets are global, the markets upon which they trade are often regional and have varying attributes. As Benedetti and Nikbakht (2021) point out, crypto-marketplaces have substantial flexibility regarding trading market design, hence have evolved into a wide variety of differentiated service offerings, specific to the target token/crypto issuers and user/investor segments catered by each marketplace. This strategic market segmentation, along with regulatory restrictions, capital controls, and the low latency in blockchain transactions, increases the cost of arbitrage within and between marketplaces, leading to persistent price disparities in cryptocurrencies.

Pieters and Vivanco (2017) study Bitcoin prices in 11 marketplaces, representing 26% of global volume, from June 2014 to July 2015 and find significant price disparities and relate them to the characteristics of each marketplace. The main differentiating attribute of marketplaces with the largest price disparities is the lack of identification requirements of users on the platform.

Choi et al. (2018) study how Bitcoin market microstructure and capital controls affect price disparities between the US and Korea. The authors identify a “Kimchi premium,” with Bitcoin trading at an average 4.73% higher price in Korea, between January 2016 and February 2018. This premium is positively correlated with Bitcoin blockchain transaction cost, confirmation time in the blockchain, and bitcoin price volatility.

Makarov and Schoar (2019) and Borri and Shakhnov (2018) find large price disparities for bitcoin, ether, and XRP in marketplaces within the same country and even larger between countries and geographic regions. These disparities are persistent and asymmetric in nature, leading bitcoin to usually trade at a premium relative to the price in US marketplaces. For instance, in Korea from December 2017 to January 2018, bitcoin traded at an average 40% premium compared to the US. For the same period, the average prices in Japan and Europe were 10% and 3% higher than in the US, respectively. Makarov and Schoar find evidence of the arbitrage spread being positively correlated with capital controls and expanding during periods of rapid price appreciation.

5.5 Regulatory Constraints

The scope of financial applications affected by blockchain technology is growing constantly. This leads to an ever expanding set of applicable regulations to blockchain developments, and also challenges current regulations to continuously adjust to the scope of influence of blockchains in the financial ecosystem.

Regulating within the open, pseudonymous, and distributed nature of permissionless blockchains is substantially more challenging than regulating entities with more traditional legal structures such as corporations. It is comparable to the efforts undertaken in the past decades with the goal of regulating the Internet, but now with value changing hands.

Ducuing (2019) presents key regulatory challenges: identifying, defining, classifying, and regulating new technologies and applications quickly enough to keep up with the current speed of developments; identifying the competent institution, government or regulatory body whom should bring forward such

regulation and lastly; implementing regulation in a way that promotes, or at least does not hinder, blockchain innovation.

As noted by Paech (2016), anonymity and pseudonymity of current public blockchains have led to the use of cryptocurrencies for illegal purposes. Moreover, even if not anonymous, the open and distributed nature prevents the implementation of common know your customer and anti-money-laundering regulations.

The possibility to create smart contracts, self-executing transactions according to a predefined set of rules, is usually touted as one of the key disruptions of blockchain technology. This disruption is perhaps one of the key regulatory challenges. As mentioned in Paech (2016), while ex-ante regulatory measures might decrease the risk of smart contract loopholes; ex-post, once these smart contracts are executed, the possibility of recourse, annulment or simply correcting the outcome of transactions will depend not only on applicable regulation, but the technological capabilities of performing such actions.

On a more general topic, the requirements of the General Data Protection Regulation (GDPR) are incompatible with current blockchain technology. Finck (2018) points out that GDPR was conceived and structured with the notion of centralized data collection, storage and processing, while blockchain technology aims to decentralize each process. According to GDPR, data controllers must enforce certain personal data rights and obligations. In an open and permissionless blockchain, each validating node would be considered an independent data controller, hence subject to GDPR and responsible for ensuring the rights of each user. These rights (commonly summarized as the right to amendment, to access and to be forgotten) also antagonize one of the key characteristics (immutability) of data storage on blockchain technology. Moreover, even though cryptographically secured (or “hashed”) information stored on blockchains is not technically readable/usable by unauthorized parties, it still qualifies as personal data, regulated under GDPR. The challenge, therefore, is to modify and shape regulation in the light of new technology, as opposed to modify and shape technology to current regulation.

5.6 Scaling

Many contemporary blockchains are limited in terms of the number of transactions they are able to process when compared to traditional systems. VISA is capable of processing at least 54,000 transactions per second (Vermeulen 2017), and as of 2011 NASDAQ routinely handled 500,000 messages per second with a capacity upwards of one million messages per second

(Miller 2011). Any blockchain that is tightly integrated with modern trading volumes must be able to accommodate a high message processing speed.

Blockchains have varying transaction throughput ability, depending on technical features. Proof-of-work blockchains like Bitcoin or Ethereum tend to have very limited transaction bandwidth in their unmodified states, with Bitcoin capable of processing less than 7 transactions per second (Poon and Dryja 2016). Many view the transaction capacity of these chains as a technical issue that can be solved by introducing layer two protocols (Poon and Dryja 2016) or sharding (Zamani et al. 2018).

Several other chains have been designed to address the transaction limitation. Ripple, a distributed ledger targeted at institutional asset transfers, claims to be able to scale to 65,000 transactions per second. Stellar, EOS, IOTA, and Qtum are other examples of existing chains that claim to exceed 1,000 transactions per second (O'Neal 2019).

5.7 Hard Forks

Distributed ledgers are more fluid records of transaction history compared to traditional transaction systems. Ledgers gradually finalize over time, as blocks are added to the chain and it becomes more costly to challenge the consensus. It is possible that a distributed ledger may diverge, where two separate transaction histories or two ledgers with different rules may be generated. This process is called a fork, and they are an important feature of any blockchain system. Forks come in multiple types. Interested readers are referred to Catalini and Gans (2016) or Biais et al. (2019) for a more in-depth taxonomy. For the moment we will focus on hard forks, and their impact on trading systems.

In a hard fork, the rules by which a ledger operates are changed, or the history of transactions can be revised (Abadi and Brunnermeier 2018). After a fork occurs, there are two branches of the ledger, one of which is the original ledger, and the other is a copy of the original with some modifications. Both typically share the same history leading up to the fork, and subsequently diverge following the fork.

After a fork, validators must make a choice regarding which branch they support. When the forking event is a simple bug fix, this is a simple decision. Validators should simply adopt the chain with unambiguously less faulty code. But the decision is significantly less clear when the rules of the system change or the record of transactions differ. Some new branches will continue alongside the original without disappearing. Biais et al. (2019) demonstrate

analytically how this can occur, as participants separate themselves onto separate chains.

Many view hard forks as a negative aspect of blockchain-based systems, but Abadi and Brunnermeier (2018) view them as a useful feature that ensures the resilience of any blockchain system. In their model, a permissionless blockchain reduces the economic rents that centralized record keepers like custodians and exchanges take. Importantly, public blockchains allow for the portability of information, such that anyone who wishes to improve a system may simply copy the data on the original chain, but use a new feature or set of rules going forward.

Abadi and Brunnermeier (2018) are markedly less concerned with attacks on blockchains from an equilibrium perspective. Hard forks can be used to circumvent an attack *ex-post*, since legitimate users can agree to fork the chain and create a history where an attack never occurred, thus reducing the risk considerations highlighted in Walch (2015) or Bonneau (2019).

A fork generates a second version of an asset. In the case of cryptocurrencies, the value is determined by the market and it's possible they both find an audience of market participants that assign value to them. However, in the case of tokens that represent ownership claims on real world assets, hard forks present a unique challenge since only one can be considered the true ownership claim. For example, if a corporation issues shares of stock as tokens on the Ethereum blockchain, and the Ethereum blockchain is subsequently forked, then two versions of the token now exist. It will be incumbent upon issuers and issuance platforms, and perhaps the legal system, to address which chain contains the valid claim on the asset in these cases.

5.8 Undesired Transparency

Miraz and Donald (2018) point out that major securities markets are relatively dark today, due to fragmented pricing in segmented markets, as well as depository or “street name” consolidated ownership that obscures the true owner’s identity. Trading assets on blockchains offers a path toward multilateral transparency.

While transparency is often described as a feature of public blockchains, many financial market participants view it as a bug. When investment strategies are proprietary, adopting a system that offers visibility into holdings and trading patterns is a non-starter for many asset managers. For this reason, various technological privacy solutions have recently come to the fore to simultaneously enable both provably correct state of the blockchain ledger

and obscured transaction activity. Examples include zero knowledge proofs and ring signatures.

6 Conclusion

The term “blockchain trading” encompasses a wide variety of market activity ranging from spot markets for cryptocurrency, to back end infrastructure on traditional asset exchanges. As highlighted in this chapter, the spectrum of trading venues, asset types, and service providers is vast. Additionally, new applications are arriving even beyond those that we address, such as the market for energy trading. We review the extant literature on various benefits and challenges to widespread adoption of blockchains in trading and exchange, but note that there will be numerous avenues for future research by financial scholars as these systems continue to develop and evolve.

References

- Abadi, Joseph, and Markus Konrad Brunnermeier. 2018, December. Blockchain economics. CEPR Discussion Paper No. DP13420. Available at SSRN: <https://ssrn.com/abstract=3310346>.
- Alexander, Carol, Jaehyuk Choi, Heungju Park, and Sungbin Sohn. 2019, March 15. BitMEX bitcoin derivatives: Price discovery, informational efficiency and hedging effectiveness. Available at SSRN: <https://ssrn.com/abstract=3353583> or <https://dx.doi.org/10.2139/ssrn.3353583>.
- Ante, Lennart. 2019, September 9. Market reaction to exchange listings of cryptocurrencies. Available at SSRN: <https://ssrn.com/abstract=3450301> or <https://dx.doi.org/10.2139/ssrn.3450301>.
- Aune, Rune, Maureen O'Hara, and Ouziel Slama. 2017, January 10. Footprints on the blockchain: Information leakage in distributed ledgers. Available at SSRN: <https://ssrn.com/abstract=2896803> or <https://dx.doi.org/10.2139/ssrn.2896803>.
- Benedetti, Hugo, Nikbakht, Ehsan. 2021, February. Returns and network growth of digital tokens after cross-listings. *Journal of Corporate Finance* (66): 101853, doi.org/10.1016/j.jcorpfin.2020.101853.
- Benos, Evangelos, Rod Garratt, and Pedro Gurrola-Perez. 2017, August 22. The economics of distributed ledger technology for securities settlement. Available at SSRN: <https://ssrn.com/abstract=3023779> or <https://dx.doi.org/10.2139/ssrn.3023779>.
- Biais, Bruno, Christophe Bisière, Matthieu Bouvard, Catherine Casamatta. 2019, May. The blockchain folk theorem. *The Review of Financial Studies* 32 (5): 1662–1715, <https://doi.org/10.1093/rfs/hhy095>.

- Bonneau, J. 2019, February. Hostile blockchain takeovers (short paper). In *International Conference on Financial Cryptography and Data Security*, pp. 92–100. Springer, Berlin, Heidelberg.
- Borri, N., and K. Shakhnov. 2018. Cryptomarket discounts. Available at SSRN 3124394.
- Bramhall, Park. 2019, June 19. Blockchain isn't always the solution (or why tokenizing equity securities is not the answer to the proxy voting problem). Available at SSRN: <https://ssrn.com/abstract=3393774> or <https://dx.doi.org/10.2139/ssrn.3393774>.
- Budish, E., P. Cramton, and J. Shim. 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130: 1547–1621. <https://doi.org/10.1093/qje/qjv027>.
- Buterin, V. 2016, September 22. Transaction spam attack: Next steps. Retrieved September 24, 2019, from <https://blog.ethereum.org/2016/09/22/transaction-spam-attack-next-steps/>.
- Catalini, C., and J. S. Gans. 2016. Some simple economics of the blockchain. Working Paper No. 22952. <https://doi.org/10.3386/w22952>.
- Caytas, Joanna. 2016, June 22. Developing blockchain real-time clearing and settlement in the EU, U.S., and globally. *Columbia Journal of European Law: Preliminary Reference*. Available at SSRN: <https://ssrn.com/abstract=2807675>.
- Chohan, Usman W. 2017, December 8. Blockchain and securities exchanges: Australian case study. Available at SSRN: <https://ssrn.com/abstract=3085631> or <https://dx.doi.org/10.2139/ssrn.3085631>.
- Chiu, Jonathan, and Thorsten V. Koepl. 2019, May. Blockchain-based settlement for asset trading. *The Review of Financial Studies* 32 (5): 1716–1753. <https://doi.org/10.1093/rfs/hhy122>.
- Choi, K. J., A. Lehar, and R. Stauffer. 2018. Bitcoin microstructure and the Kimchi premium. Available at SSRN 3189051.
- Daian, P., S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, A. Juels, et al. 2019. Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. arXiv preprint [arXiv:1904.05234](https://arxiv.org/abs/1904.05234).
- Ducuing, C. 2019. How to make sure my cryptokitties are here forever? The complementary roles of blockchain and the law to bring trust. *European Journal of Risk Regulation* 10 (2): 315–329.
- Eskandari, S., S. Moosavi, and J. Clark. 2019. SoK: Transparent dishonesty: front-running attacks on Blockchain.
- Easley, David, Maureen O'Hara, and Soumya Basu. 2019. From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics* 134: 91–109.
- Finck, M. 2018. Blockchains and data protection in the European Union. *European Data Protection Law Review (EDPL)* 4 (1): 17–35.
- Fox, S. 2017. Recent cases demonstrate need for blockchain. Harvard Law School Forum on Corporate Governance and Financial Regulation.

- Griffin, John M., and Amin Shams. 2019, October 28. Is bitcoin really un-tethered? Available at SSRN: <https://ssrn.com/abstract=3195066> or <https://dx.doi.org/10.2139/ssrn.3195066>.
- Haber, S., and W.S. Stornetta. 1991. How to time-stamp a digital document. *Journal of Cryptology* 3: 99–111. <https://doi.org/10.1007/BF00196791>.
- Johnson, B., A. Laszka, J. Grossklags, M. Vasek, and T. Moore. 2014. Game-theoretic analysis of DDoS attacks against bitcoin mining pools. In *Financial cryptography and data security*, ed. R. Böhme, M. Brenner, T. Moore, and M. Smith, 72–86. Berlin Heidelberg: Springer.
- Kahan, M., and E. Rock. 2007. The hanging chads of corporate voting. *Georgetown Law Journal* 96: 1227.
- Khapko, M., and M. Zoican. 2019. How fast should trades settle? (SSRN Scholarly Paper No. ID 2881331). Retrieved from Social Science Research Network website: <https://papers.ssrn.com/abstract=2881331>.
- Kroll, J.A., I.C. Davey, and E.W. Felten. 2013. The economics of bitcoin mining, or bitcoin in the presence of adversaries, 21.
- Mainelli, Michael, and Alistair K.L. Milne. 2016, May 9. The impact and potential of blockchain on the securities transaction lifecycle. SWIFT Institute Working Paper No. 2015–007. Available at SSRN: <https://ssrn.com/abstract=2777404>.
- Malinova, Katya, and Andreas Park. 2017, July 26. Market design with blockchain technology. Available at SSRN: <https://ssrn.com/abstract=2785626> or <https://dx.doi.org/10.2139/ssrn.2785626>.
- Makarov, I., and A. Schoar. 2019. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*.
- McCorry, P., S. F. Shahandashti, and F. Hao. 2017, April. A smart contract for boardroom voting with maximum voter privacy. In *International Conference on Financial Cryptography and Data Security*, 357–375. Cham: Springer.
- McPartland, K. 2016. Corporate bond electronic trading continues growth trend. Tech. rep., Greenwich Associates.
- Micheler, Eva, and von der Heyde, Holding Luke. 2016, May 31. Clearing and settling securities through blockchain technology creating an efficient system by empowering asset owners. Available at SSRN: <https://ssrn.com/abstract=2786972> or <https://dx.doi.org/10.2139/ssrn.2786972>.
- Miller, R. 2011, March 15. NASDAQ steps up its game in the data center. Retrieved September 22, 2019, from Data Center Knowledge website: <https://www.datacenterknowledge.com/archives/2011/03/15/nasdaq-steps-up-its-game-in-the-data-center>.
- Mills, David, Kathy Wang, Brendan Malone, Anjana Ravi, Jeff Marquardt, Clinton Chen, Anton Badev, Timothy Brezinski, Linda Fahy, Kimberley Liao, Vanessa Kargenian, Max Ellithorpe, Wendy Ng, and Maria Baird. 2016. *Distributed ledger technology in payments, clearing, and settlement*. Finance and Economics Discussion Series 2016-095. Washington: Board of Governors of the Federal Reserve System. <https://doi.org/10.17016/FEDS.2016.095>.

- Miraz, M.H., and D.C. Donald. 2018, August. Application of blockchain in booking and registration systems of securities exchanges. In *2018 International Conference on Computing, Electronics & Communications Engineering (iCCECE)*, 35–40. IEEE.
- Moore, T., N. Christin, and J. Szurdi. 2018. Revisiting the risks of bitcoin currency exchange closure. *ACM Transactions on Internet Technology (TOIT)* 18 (4): 50.
- Nakamoto, S. 2008. Bitcoin: A peer-to-peer electronic cash system.
- Nagel, J. 2016. Markets Committee Electronic trading in fixed income markets. Tech. Rep. 9789291974207, Bank for International Settlements.
- Oliver Wyman/Swift. 2014. The capital markets industry: The times they are A-Changin'.
- O'Hara, M. 2015. High frequency market microstructure. *Journal of Financial Economics* 116 (2): 257–270. <https://doi.org/10.1016/j.jfineco.2015.01.003>.
- O'Neal, S. 2019, January 22. Who scales it best? Inside blockchains' ongoing transactions-per-second race. Retrieved September 22, 2019, from Cointelegraph website: <https://cointelegraph.com/news/who-scales-it-best-inside-blockchains-ongoing-transactions-per-second-race>.
- Paech, Philipp. 2016, December 16. The governance of blockchain financial networks (2017). *Modern Law Review* 80 (6): 1073–1110; LSE Legal Studies Working Paper No. 16/2017. Available at SSRN: <https://ssrn.com/abstract=2875487> or <https://dx.doi.org/10.2139/ssrn.2875487>.
- Patti, Giovanni. 2018, December 31. The Regulation of Financial Product Innovation Typified by Bitcoin-Based Derivative Contracts. *Review of Banking and Financial Law*. Forthcoming. Available at SSRN: <https://ssrn.com/abstract=3380770>.
- Peters, G.W., and G.R. Vishnia. 2018. Blockchain architectures for electronic exchange reporting requirements: EMIR, Dodd Frank, MiFID I/II, MiFIR, REMIT, Reg NMS and T2S. In *Handbook of blockchain, digital finance, and inclusion*, vol. 2, 271–329. Academic Press.
- Pieters, G., and S. Vivanco. 2017. Financial regulations and price inconsistencies across Bitcoin markets. *Information Economics and Policy* 39: 1–14.
- Pinna, Andrea, and Wiebe Ruttenberg. 2016, April 26. Distributed ledger technologies in securities post-trading revolution or evolution? ECB Occasional Paper No. 172. Available at SSRN: <https://ssrn.com/abstract=2770340>.
- Poon, J., and T. Dryja. 2016. The bitcoin lightning network: Scalable off-chain instant payments, 59.
- Vermeulen, J. 2017. VisaNet—Handling 100,000 transactions per minute. Retrieved September 22, 2019, from <https://mybroadband.co.za/news/security/190348-visanet-handling-100000-transactions-per-minute.html>.
- Walch, Angela. 2015, March 16. The bitcoin blockchain as financial market infrastructure: A consideration of operational risk. *18 NYU Journal of Legislation and Public Policy* 837 (2015). Available at SSRN: <https://ssrn.com/abstract=2579482>.

- Wright, Aaron, and Primavera De Filippi. 2015, March 10. Decentralized blockchain technology and the rise of Lex Cryptographia. Available at SSRN: <https://ssrn.com/abstract=2580664> or <https://dx.doi.org/10.2139/ssrn.2580664>.
- Zamani, M., M. Movahedi, and M. Raykova. 2018. RapidChain: Scaling blockchain via full sharding. Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security - CCS 18, 931–948. <https://doi.org/10.1145/3243734.3243853>.



Initial Coin Offerings: Current Research and Future Directions

Jiasun Li and William Mann

1 Introduction

In recent years, many new crypto assets have been created through a process known as an initial coin offering (ICO). While the name is clearly reminiscent of an initial public offering (IPO), this direct comparison is misleading, as the two processes are in fact very different: A typical IPO issuer tends to be a relatively large, mature, and widelyheld corporation, and is usually profitable or at least revenue-generating by the time of the IPO. Moreover, the IPO is associated with the sale of a specific security, common equity, which represents a claim on the firm's profits and must comply with clear securities regulations. An ICO does not necessarily feature any of these characteristics.

Instead, an ICO is typically associated with the launch of a brand new venture—sometimes within an existing company, but more often as part of an entirely new startup. A typical ICO process involves the creation and sale of a class of cryptographically protected digital records (“tokens”) implemented on a blockchain (e.g., tokens implemented on the Ethereum blockchain typically

J. Li (✉)

George Mason University, Fairfax, VA, USA

e-mail: jli29@gmu.edu

W. Mann

Emory University, Atlanta, GA, USA

e-mail: william.giles.mann@emory.edu

following the ERC-20 standard), or a contract to deliver such tokens in the future (e.g., a Simple Agreement for Future Tokens, or SAFT).

These tokens can later be used for a number of different purposes, depending on the setup of the ICO: Often their purposes are to redeem a product or service in the future, or to serve as a medium of exchange among users on a platform to be launched later. In these cases, the owners of the tokens form a community of potential users once the product or platform becomes available. Such tokens are often labeled as “utility” tokens, reflecting the fact that their value derives from their use rather than from any associated cash flow rights. Another common category includes so-called “security” tokens, which entitle their holders to some form of cash flows, such as dividends paid out of the project’s profits or a cut of the firm’s revenues.¹ Tokens may also feature voting rights or some combination of all the above features.

The proceeds from a token sale are typically immediately available to the issuer (though vesting periods are sometimes included too). When there is a product or platform to be developed, the proceeds are typically used at least in part to fund its development. As we will summarize below, research demonstrates that ICO sales are associated with high failure rates but also with high average returns. Based on these observations, the most natural analogy for an ICO is therefore not an IPO, but rather early stage firm activities such as crowdfunding, product pre-sales, angel investing, and venture capital financing.

Furthermore, ICO tokens typically remain closely tied to the product or platform that they are initially set up for.² This stands in contrast to Bitcoin, and its competitors (such as Litecoin) and derivatives (such as Tether), which explicitly seek to become widely adopted currencies. Therefore, an ICO should be seen as a way to launch a new venture, rather than as part of a broader attempt to challenge existing payment or monetary systems. That said, ICOs do share the general goal of bypassing traditional financial intermediaries, placing them within the category of *alternative finance*, and there are clearly large overlaps in both the technology and user base between the ICO world and the broader cryptocurrency community.

From an academic perspective, there are several reasons to be interested in ICOs. The most apparent reason is simply their astonishing popularity: Aggregate ICO proceeds reached \$19 billion by June 2018 (Coindesk), and ICOs have already displaced equity issuance as the primary source of capital for blockchain-based startups (CBInsights). As with many other innovations, such rapid growth provokes both questions and controversies: From

¹ The sale of tokens that represent securities is often called a security token offering (STO).

² A major exception is Ethereum. See next section.

an investor's perspective, do ICO tokens add value to a diversified portfolio as a new asset class, or are they just redundant assets that are fundamentally worthless? The question is most striking for the case of utility tokens, which carry no cash flow rights whatsoever (marking another contrast between the ICO and the IPO). Clearly, a positive price for a utility token does not represent the present value of dividends, but rather some other form of economic value. From an issuer's perspective, could ICO issuance be useful either strategically or financially, and if so, how? Perhaps most urgently, from a regulator's perspective, is an ICO purely a dodge of securities laws that should be snuffed out as aggressively as possible, or a valuable invention that deserves a looser handling, or perhaps even encouragement?

In this article, we review the current state of academic research that explores the above questions from both empirical and theoretical perspectives. Our goal is to provide a road map to help authors embarking on this new research area to quickly get up to speed on the papers currently being circulated. Given the state of this literature, a few standard disclaimers apply, with even greater force here than with the typical review article: Most of the papers that we summarize are still unpublished working papers, and thus subject to change. Moreover, there are surely many important papers in this area that have yet to be written. Therefore, this review should be seen as a snapshot or index of current research in this rapidly growing area, and as forward-looking list of useful future research questions, rather than as a backward-looking summary of a mature literature.

Section 2 briefly overviews the history and structure of the ICO process, along with facts and terminology that should be understood before describing research in this area, and an overview of the current regulatory status of ICO tokens worldwide. The reader who is already familiar with the ICO environment can skip Sect. 2.

In Sect. 3, we summarize empirical research on ICOs. The primary goal of this growing literature is to establish broad risk and return characteristics of this new investable asset class. A second major focus has been to study the sources of token price volatility, and in some cases to connect these patterns with activity on social media. Data issues have been paramount in the literature so far, and thus we begin this section with a summary of data sources and methodologies used in empirical papers.

In Sect. 4, we review theoretical analyses of the value creation of ICOs associated with utility tokens. What is the value of a novel cryptocurrency-based approach to the well-studied challenges of launching an early stage company? The practitioner community has offered a set of intuitive rationales, but much confusion still remains, which call for rigorous economic analyses.

Economists have responded by developing various models that emphasize different aspects of ICOs. Since we view the most pressing question as being how ICOs should be viewed by regulators, we conclude with a discussion of this issue.

Section 4 ends with suggestions for future research directions, and Sect. 5 concludes.

2 The ICO Process: History, Facts, and Terminology

A typical ICO process is structured as follows:

An entrepreneur explains his or her business idea in a **white paper**, which often but does not always include the business vision, technological implementation, and token distribution methods. While it is tempting to compare a white paper with a prospectus for security issuance, in our opinion the novel nature of many high-profile ICOs makes their white paper more like an academic proposal for research grants.

For many projects, funding starts with a so-called **pre-ICO** (or presale), in which a selected group of investors are given exclusive rights to purchase tokens, typically at a discounted price, before the token sale opens up to the public.

The public token sale lasts within a pre-specified campaign window, which can range from a few days to almost a year, depending on the underlying project. Within the campaign window, tokens may be sold at a fixed price set by the issuing entity; or through a pre-determined price schedule (usually escalating automatically once a certain amount of calendar or block time has elapsed, or certain number of tokens have been sold); or sometimes by a market-clearing mechanism. ICO tokens are usually purchased using ETH or BTC. Pre-sales and SAFTs sometimes also accept fiat currencies.

Many ICOs include a **hard cap**, a maximum amount of tokens to be issued during the campaign window. A hard cap rations the tokens sold in the ICO: if it is reached before the campaign window ends, the ICO automatically closes, and no further tokens will be issued. Many ICOs also include **soft caps**, which is similar to the “all-or-nothing” clause commonly observed in many crowdfunding campaigns. If the amount of funds received by the end of the campaign window does not reach the soft cap, the ICO will be canceled, and all received funds returned. Many analysts and researchers view reaching either cap as a measure of ICO “success,” although it must be kept in mind

that the choice of the cap levels is an endogenous decision of the issuer, and can intentionally be set to levels that are easily attained.

An important and distinctive feature of the ICO process is that all the above procedures (such as terminating the campaign when the window closes or when the hard cap is reached, returning funding when the soft cap is not reached, etc.) can be hard-coded in a “smart contract,” and can thus be automatically implemented without the involvement of any third-party entities. Hence, compared to traditional funding methods, an ICO has the potential to be more transparent and “trustless,” in the sense of being less dependent on trusted intermediaries. Not all ICOs take full advantage of this potential, but based on our private communications with practitioners the failure to do so is often viewed as a signal of inferior quality.

Throughout the ICO process and subsequent project development, token issuers actively advertise their ICOs and engage with potential token buyers through various channels. In addition to dedicated websites and the white papers, social media also play an important role in bringing awareness to potential token buyers. Potential ICO token buyers are actively engaged in discussions about the project on social networks such as Twitter, Telegram, Discord, Reddit, 4chan, Facebook, or WeChat.

Tokens issued in an ICO are often listed later on one or more cryptocurrency exchanges. Listing on a major exchange is often viewed as a marker of success, based on the idea that a major goal of the issuer following the ICO is to provide early stage investors with liquidity on the secondary market (again, analogous to the IPO process). However, it is not necessarily clear that token liquidity is always desirable for projects at such an early stage, and this question is a subject of ongoing research (see later discussion). Tokens may also be traded over-the-counter through so-called decentralized exchanges (e.g. EtherDelta, 0x, etc.), although these decentralized exchanges constitute a very small market share, and are currently facing legal challenges (for example, the SEC recently prosecuted the founder of EtherDelta for violating security laws).

Token buyers are geographically distributed all around the world, and in theory political borders impose no restrictions on these transactions. However, in practice it is becoming increasingly common for token issuers to exclude buyers with IP addresses from certain jurisdictions. For example, many deals bar Chinese buyers, following a ban put in place by the Chinese government in September 2017. Some deals also bar US buyers due to regulatory uncertainty, as discussed below.

Although no universally accepted definitions prevail, market participants often distinguish the terms “token” and “coin.” A “token” operates on an

existing blockchain, while a “coin” has its own blockchain. Under this definition, most “initial coin offerings” are actually sales of tokens, usually implemented by smart contracts on the Ethereum blockchain.³ The ERC-20 protocol standardizes the smart contract. Sometimes a project launching a new blockchain issues tokens on an existing blockchain first during the ICO stage, with the plan to convert the tokens into coins on the project’s native blockchain once developed. The distinction between “tokens” and “coins” is not critical to our discussion, and we do not consistently respect it in this article.

In addition to tokens sold to private investors during the pre-ICO, and to the public during the actual ICO, often token issuers also retain a fraction of the total token supply, which may be useful as future remuneration to developers, inventory for future fundraising, or as a signal of project quality (analogous to classic models of signaling via equity retention). Some projects also retain tokens as incentives for some future stakeholders that are key to the development of a user community of the project in question, for example, as “mining rewards” if a native blockchain is involved or as “verifier incentives” if some alternative consensus methods are planned (e.g. side-chain solutions that are under active research). Some projects also promise to buy back tokens either at a pre-set pace or at issuer discretion. The later sale or repurchase of tokens may have effects similar to open market operations by a central bank in currency markets, or share repurchases by a corporation in equity markets. However, these issues have not been carefully analyzed in the literature, perhaps because the platforms supported by ICOs are still too young to provide adequate empirical evidence.

Several classifications have been proposed for the purpose of ICO tokens. In the usual terminology, “utility tokens” carry the right to redeem a product/service, or to participate in a platform, at a later date. “Work tokens” belong to a related and sometimes overlapping category to utility tokens, in which tokens are used as licenses for developers to run decentralization applications on the platform. “Security tokens” carry cashflow rights, and thus are essentially tokenized securities. We note, however, that no consensus has been reached yet on the proper classification of individual tokens. Given the flexibility of token usage and rapid development of the ICO space, we expect more token models and classification schemes to emerge in the years to come, and the boundaries between these categories may become increasingly blurry.

Compared to all startups, a disproportionate number of ICO startups have venture capital backing, although VC-backed ICOs are still in the minority.

³ Howell et al. (2018) report that only 9% of ICOs in their sample planned to create their own blockchain (Table 2, Panel 1).

For example, Howell et al. (2018), discussed below, find that 11% of ICO issuers in their sample have received VC funding. VC funds are aware of the potential for token issuance by their portfolio companies, and some now demand rights to potential future token sales when they invest, which highlights an interesting capital-structure issue with token issuance. Chod and Lyandres (2018), discussed below, investigate this issue and the relationship between token sales and equity offerings.

Finally, for illustration we list a few notable ICO examples:

MasterCoin, known as Omni today, is commonly regarded as the first ICO. The goal of Mastercoin is to add new features to the core Bitcoin protocol. In order to achieve wide distribution of Mastercoins, its developers launched a crowdsale—selling Mastercoins in exchange for Bitcoins in 2013. A notable feature of Mastercoin’s crowdsale is that all the Bitcoin spent were cryptographically “burned”—sent to an unspendable address and thus permanently eliminated from circulation. This example vividly illustrates that an ICO may not necessarily have a financing purpose, and can purely serve as a way to build up a platform or community. The platform-building aspect of ICOs is modeled in Li and Mann (2018), discussed below.

Ethereum, which is commonly used as the platform for other ICOs today, originally started from a crowdsale of its native token, Ether, back in 2014. In June 2018, William Hinman, director of the SEC’s division of corporation finance, famously argued that “based on my understanding of the present state of Ether, the Ethereum network, and its decentralized structure, current offers and sales of Ether are not securities transactions,” even while acknowledging “the fundraising that accompanied the creation of Ether.”

FileCoin, while not explicitly referred to by the issuers as an ICO, is nevertheless a widely cited example for several reasons: First, its tokens are designed to serve as the internal medium of exchange within a peer-to-peer file storage platform, a business model that is representative of many high-profile ICO deals. Second, it was the largest ICO ever by the time of its issuance. Third, it popularized the Simple Agreement for Future Tokens model. For a detailed case study of the FileCoin offering, see Howell et al. (2018).

3 Empirical Research on Initial Coin Offerings

In this section, we review a selected sample of recent empirical papers on ICOs in detail. We refrain from the impossible task of fully covering every paper in this rapidly growing literature, but rather use a few papers as a

guide to understanding key issues in empirical research in this area. A longer (though still incomplete) list of papers appears at the end of this section.

It goes without saying that the empirical findings in the literature to date can only speak to ICOs from recent history, as this is the only time period available to study. With hindsight, this current era may eventually appear to have been anomalous. This issue will naturally be addressed by the passage of this time. For now, one must simply remember to be cautious about extrapolating recent empirical findings to the future.

3.1 Data Sources: Overview

There is not yet a one-stop data source for ICO sample construction. Furthermore, there are fundamental challenges in constructing a comprehensive, accurate, and unbiased ICO sample that is free of selection issues or survivorship bias. In this sense, the ICO literature mirrors the venture capital literature.

One practical issue in studying returns and proceeds to ICOs is their proper denomination. For easy interpretation in real terms, and for comparability with other literature, it would seem most natural to measure them in dollars. On the other hand, much of the variation in dollar measurements would be driven by the notoriously high volatility of Bitcoin and Ethereum prices, against which much ICO activity is benchmarked. The literature has taken various approaches to this issue, and we do not attempt to describe each one in detail, but only to highlight the issue for the reader.

Every project that we have reviewed combines data from two major sources: First, data on prices and trading volumes for ICO tokens come from cryptocurrency exchanges or exchange aggregators, most notably CoinMarketCap.com. Second, data on ICO issuers, dates, and structures come from specialized ICO listing websites. Here, we describe the leading examples in each category:

CoinMarketCap.com (CMC) is the leading aggregator of market data on cryptocurrencies, and the most popular data source for empirical work on ICOs, used in every empirical paper we discuss below. CMC allows cryptocurrency exchanges, and individual coins and tokens, to provide APIs that feed transaction data. Individual coins and tokens are also listed if CMC receives a request and determines that the token complies with listing requirements. For the coins and tokens that it tracks, CMC provides daily historical opening, closing, high, and low prices, as well as volume and market capitalization (though it may not have complete coverage of all of these entries for each coin or token). As of November 2018, there are 2088 cryptocurrency

coins and tokens being tracked. CMC earns revenue from advertising, user donations, and referral commissions, and possibly from other sources as well.

The listing requirements on CMC, which attempt to maintain the investment quality of its listed coins and tokens, have several practical implications for researchers. One is that CMC's listing standards may vary over time, and in ways that seem subjective from an outsider's perspective. For example, CMC removed some South Korean cryptocurrency exchanges when it saw their prices diverge substantially from other venues.⁴ Conversely, CMC stopped imposing a requirement of minimum trading volume in July 2018, when concerns arose that exchanges were able to game this requirement through various methods including "wash" trades.⁵

These incidents highlight that CMC data may be subject to survivorship bias, in ways that are not completely clear and that may vary over time along with the listing and delisting standards, even though CMC does keep historical records for delisted tokens, which likely ameliorates this concern.⁶ The concerns over "wash" trades also highlight that cryptocurrency transaction data must be interpreted with care. Even if reported transaction volumes aggregated by CMC are real, they do not necessarily have the intended economic interpretation. Finally, as noted by Benedetti and Kostovetsky (2018), discussed below, the review process for token listings may impose a lag in the appearance of price data on CMC.

Competitors to CMC include CoinGecko, which is used in a few of the papers we discuss below, as well as some more recent products from John McAfee and from Forbes. Less commonly, researchers may also download the data provided on these aggregators directly from various individual exchanges.

Unlike secondary market transactions that are not directly recorded on blockchain, primary market transactions of ERC-20 tokens can be obtained directly from the Ethereum blockchain, as is done in Lee et al. (2018). The clear advantage of this approach is that it addresses concerns over selection and survivorship bias created by the listing and delisting standards of individual exchanges, though with the implicit assumption that ICOs following the ERC-20 protocol are the population of interest or are representative of that population.

⁴ <https://cointelegraph.com/news/coinmarketcap-removes-south-korea-exchanges-ripple-market-cap-drops-20-billion>.

⁵ <https://blog.coinmarketcap.com/2018/07/19/coinmarketcap-stands-for-data-transparency-and-clarity/>.

⁶ In an interview from 2015 (when volume requirements were still in place), a representative from CMC suggested that 40% of tokens ever added to the website to that point had been delisted. Tokens have also been delisted when the API of the associated exchange stopped working properly. <https://cointelegraph.com/news/coinmarketcap-about-40-of-the-coins-ever-added-to-the-site-are-now-inactive>.

ICO listing websites charge fees to list ICOs, with the target audience being potential investors. The websites typically offer information about the ICO including start and end date, the percentage of tokens ultimately sold and the capital raised, the price schedule during the ICO, any hard or soft caps, and a dummy for whether there was a pre-ICO round; and information on the issuer including website, social media handles, and the country of registration. The most common ICO listing websites used in empirical research that we have reviewed to date are *tokendata.io*, *icobench.com*, and *coinschedule.com*, but several more have also been used and are also listed below.

Several potential sources of selection bias arise with these websites. Given the pay-to-play model, there is no guarantee that the tokens listed on any individual website are representative of the universe of investable tokens. Individual websites may have different approaches to backfill information on ICOs launched before the website existed. Perhaps most importantly, survivorship bias can severely affect the computation of unconditional risk and return characteristics. Aside from selection issues, there may also be inaccuracies in data or inconsistencies in definitions. These issues are the subject of ongoing investigation in empirical work, as we describe below (Table 1).

3.2 Empirical Findings from Selected Papers

We first summarize several descriptive findings and regularities in the ICO space. We draw on statistics from the first two papers that we discuss in detail below: Howell et al. (2018), and Benedetti and Kostovetsky (2018). However, the patterns described here are broadly similar across all the empirical literature to date.

Trends in ICO activity: ICO activity exploded in the second quarter of 2017, before stabilizing to lower volumes thereafter. Figure 1 of Howell et al. (2018) shows that aggregate proceeds in their sample peaked at \$6bn in this quarter, then hovered around \$3bn in subsequent quarters. Figure 2 of Benedetti and Kostovetsky (2018) shows that the number of ICOs increased throughout this period, but the increase is driven mainly by ICOs that raised relatively little capital and ultimately did not list on exchanges.

Geography: Figure 1 in Benedetti and Kostovetsky (2018), and Figure 2 in Howell et al. (2018), show that the United States is by far the country with the greatest number of ICO issuers in absolute terms, followed by Russia, China, India, and (as a group) western European countries.

Table 1 Below lists several empirical papers and the datasets they use

	Coinmarketcap.com (CMC)	Tokendata.io	Icobench.com	Icorating.com	Smithandcrown.com	coindesk.com	coinschedule.com	CMC	Additional sources
AGM	X	X		X	X		X	253	lcoalert.com, icobazaar.com, tokenmarket.net
AS	X	X	X			X	X	363	Cryptoslate.com, coinmarketplus.com
BDEM	X	X	X	X	X		X	659	lodata.io,
BK	X	X	X	X				609	ico-check.com, icodrops.com, bitcointalk.org
DLZ	X					X			CoinGecko, CryptoCompare
Fisch	x						X		icodrops.com, tokenmarket.net
HNY	X	X						453	Proprietary data from FileCoin ICO
HPR	X							64	222 digital coins; hand-collected large ICOs

(continued)

Table 1 (continued)

	Coinmarketcap.com (CMC)	Tokenata.io	Icobench.com	Icorating.com	Smithandcrown.com	coindesk.com	coinschedule.com	Sample size, after matching to	Additional sources
LLS	X	X	X	X	X			433	Primary market transactions from the Ethereum blockchain
LY	X							893	Secondary market including non-ICO assets
Momtaz	X							302	

Note: AGM = Adhami, Giudici, and Martinazzi (2018). AS = Amsden and Schweizer (2018). BDEM = Bourveau et al. (2018). BK = Benedetti and Kostovetsky (2018). DLZ = Deng, Lee, and Zhong (2018). Fisch = Fisch (2018). HNY = Howell, Niessner, and Yermack (2018). HPR = Hu, Parlour, and Rajan (2018), LLS = Lee, Li, and Shin (2018). LY = Li and Yi (2018), Momtaz = Momtaz (2018)

Structural details: The mean ICO duration was 37 days in Benedetti and Kostovetsky (BK), and 40 days in Howell et al. (HNY). A presale (or pre-ICO round) was conducted by 40% of issuers in BK, and 45% in HNY. The mean percentage of total tokens sold in the ICO was 60% in BK, and 54% in HNY. A fundraising goal was reported by 61% of issuers in the sample of HNY, and of these, 47% achieved it. In the sample of BK, 26% of tokens had been listed on an exchange by the end of their sample.

Skewness in proceeds: The mean (median) dollar amount an ICO raises is \$15.8 million (\$6.62 million) in the sample of Howell et al. (2018), and \$11.5 million (\$3.8 million) in Benedetti and Kostovetsky (2018). These numbers should not be directly compared across papers, as they are sensitive to the details of sample construction in a given paper (see below); but both reflect a heavily skewed distribution of ICO proceeds, with a small number of highly successful transactions driving much of the aggregate activity.

Skewness in post-ICO token price appreciation: This is one of the most important outcomes to study from an economic perspective, but also one of the most difficult, again due to extreme skewness. In Benedetti and Kostovetsky (2018), listing ICOs exhibit a median return of 21%, but a mean return of 246%, from the ICO closing price to the first opening price.⁷ The apparent short-run underpricing of ICOs has motivated several empirical papers, by an analogy to the IPO market. Over a longer horizon, Howell et al. (2018) find that, conditional on listing on CMC, the mean five-month cumulative excess return over Bitcoin is +150%, while the median is -50%.

Based on the last two stylized facts, the proceeds and investment return of “typical” ICOs are difficult to measure or even to define: A few ICOs have been phenomenally successful, many more only modestly so, and a large number have generated no value or even significant losses. Of course, this is characteristic of any skewed distribution, and is nothing abnormal in the world of startup financing. Venture capital firms, for example, make most of their profits from only a few investments, while the dotcom bubble generated many failed companies but also gave rise to success stories like eBay and Amazon.

Having presented some general findings, we next discuss a few empirical papers in detail, highlighting specific contributions both in terms of their analysis, and in terms of their data collection and sample construction:

Howell et al. (2018) focus on empirical predictors of liquidity in the secondary market for ICO tokens, where liquidity is measured analogously to Amihud (2002) and Amihud et al. (2006). An attractive aspect of the

⁷ These are simply summary statistics conditional on listing; the authors then take care to compute numbers that deal with selection issues, as we describe below.

paper's focus on token liquidity is that, when measured in log terms, it does not appear to be nearly as skewed as proceeds or returns. This encouragingly suggests that their empirical results capture broad patterns in the data. In contrast, a study of predictors of ICO proceeds or returns is heavily influenced by the very small number of ICOs that have achieved enormous success.

The authors find several broad predictors of ICO token liquidity. The first category includes measures of certification (prior VC backing), professional background (e.g. computer science/entrepreneurial experience by the founder), transparency (availability of a white paper and/or code on Github, and a budget for the use of proceeds), and incentive alignment (e.g. a token vesting schedule for insiders). This list overlaps with predictors of other definitions of ICO success, as investigated in other papers and described below. These general findings suggest that ICO purchasers are not completely unsophisticated, as they seem to respond to observable factors associated with project quality.

In a similar vein, tokens are more liquid when the ICO stated a fundraising goal *and* reached it, but not when a stated goal was not reached, suggesting a signaling benefit from being able to meet stated targets. (This result holds even when controlling for the amount raised, which positively predicts liquidity and is mechanically linked to the achievement of the fundraising goal.) Other structural traits of the ICO that are associated with greater future liquidity are the presence of pre-sales, acceptance of Bitcoin or Ether as payment, and implementation on the Ethereum blockchain, as well as measures of social media presence (number of Twitter followers or size of Telegram group).

Finally, with regard to the purpose of the token and the startup, utility tokens are much more liquid than other types of token. When splitting the sample across twelve different business categories, startups focused on payment systems and blockchain protocols are substantially more liquid than a reference category. These patterns do not necessarily imply that the tokens in question are inherently better, as the optimal level of liquidity may be different for different projects. (In particular, optimal liquidity is likely to be higher for utility tokens and decentralized networks, since these rely on user interactions by design.) However, these findings do help characterize the types of projects that choose the ICO mechanism, and what structural features they will choose within that mechanism.

The paper also includes a unique case study of the launch of FileCoin by Protocol Labs, aided by proprietary data from the issuer. Although FileCoin did not label its token sale as an ICO, its structure nonetheless illustrates the

“utility token” setup very clearly: The project plans to develop a decentralized platform on which users can rent disk storage space to each other, competing with providers like Amazon and Google. Using this data, the authors characterize in detail the payment methods, vesting schedules, and price evolution of the FileCoin launch.

Benedetti and Kostovetsky (2018) focus on the returns to ICO investment. A key conclusion of their analysis is that ICO tokens appear underpriced prior to initial trade on an exchange, in the sense that average returns are very high between these two dates. This stylized fact marks a connection with the IPO literature, and is consistent with evidence in other papers discussed below.

Studying returns requires careful analysis to deal with selection issues: Conditional on listing, the authors find that token prices increase by an (equal-weighted, dollar-indexed) average of 179% from the end of the ICO to the beginning of public trading, a time frame that covers just 16 days on average. However, a large number of ICO tokens do not list, so that this return is not investable unless one can identify the listing tokens *ex ante*. This motivates the authors to broaden the sample to include at least some non-listing ICOs.

To construct such a sample, the authors refer to five of the major ICO listing sites described above. In the main analysis, they consider ICOs from these websites that actually reported raising capital. Out of this broadened sample, the authors find that 75% of tokens did not list within sixty days of the end of the ICO. The unconditional returns to ICO investment thus depend critically on how one handles these ICOs, and in particular, on what return is imputed to their tokens after the end of the ICO.

The authors consider imputing -50% and -100% returns to all tokens that fail to list on an exchange. The latter is an especially aggressive approach, as at least some of the non-listing ICOs would have continued to trade off of exchanges, or else would have returned some or all of the investment funds to investors. Even in the most conservative scenario of imputing -100% returns, the equal-weighted return from the end of the ICO to the opening trade on the exchange is 31%, comparable to the typical magnitude of underpricing found in the IPO literature. The authors also test varying assumptions about the exact timing at which a hypothetical investor’s sell order would be filled during the first day (since all investors cannot sell at the open price), and find that this does not greatly impact the magnitude of underpricing.

These patterns bear a clear resemblance to the IPO process, which is also characterized by high average returns from the rationed book-building stage to the opening of unrationed trade. However, unlike with an IPO, there are

no underwriters, nor a well-defined community of institutional investors, who clearly benefit from ICO underpricing. Thus, the appearance of such a similar phenomenon in two different institutional settings may shed some light on its underlying causes in both cases. Of course, it is still theoretically possible that the returns documented here are fair compensation for the risk of ICO non-listing, which happens with higher probability than IPO withdrawal, and is more likely to result in the loss of funds.

Notably for other empirical researchers in this area, the authors also determine a hierarchy of accuracy for the data on the ICO listing websites that they use by carefully cross-checking much of their data. They arrive at the following hierarchy: (1) ICOdata, (2) ICODrops, (3) ICObench, (4) ICOrating, (5) Icocheck, with the exception that (1) and (2) are switched for price data. The paper also shows that Twitter follower count and activity are good proxies for firm size and survival, respectively, and relate these variables to post-listing ICO returns.

Lee et al. (2018), also study underpricing, as well as the predictors of ICO success, where success is defined as achieving a stated fundraising goal, or (if there is no stated goal) raising at least \$500,000. These two areas of focus are shared with, respectively, Benedetti and Kostovetsky (2018), discussed above; and Adhami et al. (2018), discussed below. A unique aspect of the analysis in Lee et al. (2018) is a focus on the information content in analyst ratings provided by ICO listing websites: In their main results, a higher analyst rating, and larger analyst following, are strongly associated with a greater probability of soft-cap success. The analyst ratings thus appear to have information content. A deeper implication is that the ICO listing websites appear to be making a meaningful effort to help purchasers discriminate between more- and less-valuable ICOs, similar to the efforts of CMC to avoid reporting data on low-quality tokens.

In other analysis, the authors find that success is more likely when a smaller fraction of tokens is offered for sale during an ICO, which is consistent with the Leland and Pyle (1977) retention signaling studied in many other fundraising contexts. Also, the authors find that ICO success is more likely if the ICO includes a presale round, which is consistent with Adhami et al. (2018), discussed below. The authors also investigate feedback effects between the decisions of individual ICO token purchasers, which is a major focus of several models discussed in the next section.

A unique aspect of this paper is that the authors manually retrieve data on token transactions from the Ethereum blockchain. This allows them to study the transactions that occur during and before the ICO process, not just if and when the token lists on an exchange. The authors use this data to analyze the

returns to early ICO investors, and also to document strong persistence in the success of the ICO from one round to the next. Also useful for other researchers is that the authors check the quality of CMC transaction data against several other exchanges, and conclude that CMC is generally reliable.

The studies described above are all still in working-paper form. As of yet, there seem to be only two empirical papers on ICOs that have been published in peer-reviewed journals:

One is **Adhami et al. (2018)**. The authors study the determinants of ICO success. As mentioned above, this focus is shared with Lee et al. (2018), though the definition of “success” is subtly different: In Adhami et al. (2018), the classification of “failed” ICOs is taken from the ICO listing websites that are their primary data source, while successful ICOs are those that are not listed as failures by these websites. Most commonly, failure on these websites refers to the failure to meet the stated soft cap fundraising goal, which is the basic definition used by Lee et al. (2018). However, failure may also include other issues like security breaches or technical failures during the ICO process.

Their logit analysis finds several predictors of ICO success. The availability of project code significantly predicts the probability of success (although interestingly the same is not true for the availability of a white paper). The disclosure of a legal jurisdiction for the project is also a strong predictor, which the authors interpret as facilitating potential legal action against the project’s founders. Finally, the existence of a presale round in the ICO is also a significant predictor.

Several of these predictors are consistent with Lee et al. (2018), and also with other papers’ predictors of “success” defined as proceeds raised or later token liquidity. Some of these robust predictors, such as the importance of disclosure through project code, are natural and reassuring evidence of at least some degree of rationality in the ICO market. Others, like the existence of a presale round, are less obvious from a theoretical perspective, which provides motivating questions for the theoretical models to be discussed in the next section.

The other published paper is **Fisch (2018)**. This paper studies the determinants of ICO proceeds. The most important predictors of a greater ICO capital raise are the length of the associated white paper and dummies for being located in the United States, and for being based on the Ethereum blockchain. Also interesting are several factors that do not seem associated with ICO proceeds, such as the price of Bitcoin, and the pre-sale indicator shown in other papers to be an important predictor of other definitions of success. As always, a fundamental feature of the distribution of proceeds is its

extreme skewness (even in log terms), meaning that the results are weighted toward a small number of highly successful ICOs.

As mentioned above, we cannot discuss in detail every paper in this area, but it is worth briefly mentioning several more. **Amsden and Schweizer (2018)** study the predictors of whether an ICO is ultimately listed on an exchange. They find that exchange listing is positively predicted by, among other things, the availability of Github code, the existence of a Telegram group, and the retention of a larger share of ICO tokens by the project insiders. We note that several of these overlap with the predictors of other measures of success in the papers discussed above, marking a reassuring consistency. **Deng et al. (2018)** examine how measures of disclosure and governance quality predict future technological development of the project. Technological development is measured using GitHub code indicators. **Bourveau et al. (2018)** study the predictors of ICO proceeds, returns, and crash risk, and **Momtaz (2018)** studies the determinants of first-day returns, which relates to the question of underpricing as articulated in other papers above. **Hu et al. (2018)** provide some investment characteristics of 64 ICOs. **Li and Yi (2018)** probe potential factor structures in the cross-sectional expected returns of crypto asset returns.

3.3 Directions for Future Empirical Research

Certainly, the major outstanding question for empirical work on ICOs is the long-term outlook for this market. The incredible growth of recent years could just as easily be followed by continued growth, stagnation, or collapse. However, this question can only be answered with the passage of time. The same is true for many important questions about the determinants and potential factor structure of long-run ICO token price performance.

In the meantime, the most interesting questions for ICO research are likely to deal with the initial demand for tokens, and how this is affected by the structuring and initial pricing of the ICO. For example, the apparent underpricing of ICO tokens outside of the traditional IPO setting is striking and potentially revealing of a deeper explanation for both patterns. Variation across ICOs in disclosure and governance standards provides an opportunity to find out which features are most important to potential investors, and pending regulatory action may provide natural experiments as a source of identification. Finally, the relationship between ICO and venture capital activity is an evolving topic, and a natural one for researchers to assess going forward.

The predictions tested in empirical research to date have mostly been generated by analogy with other financial markets like IPOs, or from simple introspection. The next step forward will almost certainly involve the testing and estimation of models that are developed specifically with ICOs in mind. To that end, we now turn to a discussion of the growing theoretical literature on ICOs that is beginning to provide those models.

4 Theoretical Research on the Value of ICOs and the Valuation of Tokens

Theoretical research has been developing rapidly along with empirical contributions, with much effort devoted to answering the simple question of why ICOs. By developing models to illustrate the economic benefits of ICOs, these papers could potentially guide regulatory decisions. Another goal of this literature is to explain pricing, deal structuring, trading volume, or other commonly observed patterns surrounding ICO tokens.

This literature has advanced several potential benefits of the ICO structure:

- **To overcome network effects:** ICOs are often associated with the launch of platforms, which typically feature network effects, defined as a strategic complementarity between the adoption decisions of individual users. A natural consequence is that there are multiple equilibria: One or more “efficient” equilibria in which platform adoption is widespread; and one or more “inefficient” equilibria in which few or no users adopt. Several papers study the problems and opportunities presented by ICOs when viewed through this lens.⁸
 - Li and Mann (2018) observe that the use of tokens can select the efficient equilibrium, since the purchase of a token credibly communicates the future intent to use the platform. Viewed this way, an ICO exploits forward induction reasoning to overcome coordination problems (as in the experiment of van Huyck, Battalio, and Beil, 1993, among others). This argument captures two important features of platform-specific ICO tokens: First, the aggregate issuance amount is perfectly observable thanks to the transparency of the smart contract/blockchain infrastructure supporting the token sale; and second, tokens, while being valuable

⁸ Ryan Zurrer, Principal & Venture Partner of Polychain Capital, has stated that “tokens act like rocket fuel for network effects.”

as digital records, are intrinsically worthless outside of the platform, at least at the time when they are first issued.

- Cong et al. (2018) analyze the interaction between token adoption and price dynamics. They assume that the surplus from holding tokens increases in the size of the user base (even though confounded by random fluctuations of platform quality over time), and focus on the nontrivial dynamic equilibrium. The speculative motives of some token holders, along with the network effect in platform value, generate multiple feedback loops that give birth to nonlinear relationships between platform quality, user base, and token price. Under certain conditions, the presence of tokens can accelerate platform adoption and stabilize the user base. However, token price volatility can be high due to both technology fluctuations and endogenous user adoption.
- Sockin and Xiong (2018) model cryptocurrency as a membership, i.e. the right to trade on a platform, within a two-period model. Network effects are explicitly specified within users' (or households' using their terminology) preferences. They highlight the presence of multiple equilibria in such settings, and derive closed-form rational expectation equilibrium that characterizes price and quantity.
- Bakos and Halaburda (2018) propose that the platform's potential users will naturally be biased towards the efficient equilibrium due to its greater salience or "focality" (Schelling, 1960). In their model, launching a platform with tokens results in lower total profits than launching without tokens, but launching without tokens requires a subsidy to early adopters. For a capital-constrained entrepreneur, subsidies may be infeasible, and so a token-based platform launch can be the optimal decision.
- **Governance:** The standard governance features of securities like equity are often lacking in ICOs. On the other hand, when they are present, they may be particularly effective thanks to the commitment technology afforded by the underlying blockchain.⁹ Therefore, governance issues can be an important consideration in adopting the ICO structure.
- Chod and Lyandres (2018) model ICOs as granting no control rights to purchasers, and thus providing a way to un-bundle cash flow from control rights. This naturally leads to a tradeoff between risk-sharing and incentives in deciding whether to adopt the ICO structure. As a result, ICOs are an attractive financing method when VCs are imperfectly diversified, when information asymmetries are not too severe,

⁹ For discussions of this point, see Yermack (2017) and Cong and He (2018).

- and/or when idiosyncratic risk is large (for example due to a skewed payoff distribution, as is commonly assumed in the startup world).
- Canidio (2018) studies the interaction of seignorage ability, incentive problems, and token prices. Seignorage—the ability to dilute token holders by issuing more—poses a governance problem that is unique to ICOs.¹⁰ Smart contracts allow the entrepreneur to commit to a deterministic token supply, but some flexibility in token supply might be optimal. In the model of Canidio (2018), the entrepreneur holds a share of the token stock to provide incentives for effort. When financing is required, the entrepreneur can sell out of this share, or (equivalently) can engage in seignorage sale of new tokens. The sale dilutes the entrepreneur’s incentives, hurting the token price; in turn, the lower price requires a greater sale. This feedback effect is a unique source of multiple equilibria compared to the papers discussed above, creating a disadvantage to using the ICO structure.
 - Garratt and Oordt (2019) model tokens as the exclusive medium of exchange for a project’s future products or services. The author compares ICO financing with debt and VC financing, and argues that ICOs could complement traditional financing methods. Specifically, for projects whose operating cost could be (probabilistically) lowered by some unobservable effort from the entrepreneur, textbook corporate finance teaches us that both debt and equity could be inefficient, because the entrepreneur bears all the cost of providing effort, but has to share the benefit of reduced cost with outside investors (as a share of the profits as in the case of equity and lowered default likelihood as in debt). To best incentivize the entrepreneur’s effort, outside investors’ payoffs should be as much separated from the benefit from any reduction in operating cost. An ICO could achieve this goal as long as the entrepreneur does not abandon the project, as in an ICO the outside investors get repaid from the revenues of selling the tokens they acquired before, which is independent of the project’s operating costs.
 - **To aggregate the “wisdom of the crowd”:** As with any pre-sale, the ICO elicits information (revealed preferences) from early adopters. This information can be useful in several ways: To the entrepreneur in deciding whether to continue the launch; to investors in setting financing terms; and to other customers in deciding whether to buy in at later dates. All of the above can contribute to the social value of the platform, and to the entrepreneur’s ability to extract some of that value through a higher price

¹⁰ In contrast, dilution does not happen with equity issuance as long as the equity is fairly priced.

during the token sale. While not their main focus, Li and Mann (2018) explicitly derive the role of ICOs in “harnessing wisdom of the crowd” in their appendix. This line of thinking closely relates to an emerging literature on crowdfunding, which is outside of the scope of this review. We direct interested readers to the references in Li (2018).

This literature has also pointed out several limitations of the ICO structure. In particular, Catalini and Gans (2018) consider a type of project where ICOs’ value creation is dubious. Their model is motivated from the observation that in many ICO cases the entrepreneur plans to develop a product that will be sold directly to consumers. Tokens in their models are essentially pre-sold rights (as well as the exclusive payment media) to purchase the products once they are developed. In other words, ICOs effectively impose a “token-in-advance” constraint for product purchase decisions. The authors show that in their model, while equity financing preserves first-best outcomes, ICOs may suffer deadweight costs. Their result cautions that a large group of common observed ICOs may indeed be misguided.

We note that many of the theoretical mechanisms described above echo results from other fields. For example, the trade of fundamentally worthless “utility” tokens at a positive price resembles money models like Kiyotaki and Wright (1989) and Townsend (1980), as well as the chip-strategy equilibria in repeated games studied by Olszewski and Safronov (2018) and prior papers. Other connections will be immediately apparent to theorists working in various subfields. While an extensive literature review is beyond the scope of this paper, we feel that these connections are part of the reason ICOs have quickly proven to be such an intriguing topic for applied theorists. It remains to be seen whether the ICOs are primarily implementations of existing mechanisms, or are completely new solutions. This observation reflects an uncertainty that applies to the ICO and blockchain universe more generally.

4.1 Regulatory Implications

Lurking behind all of the mechanisms described above is a simpler one: the desire for fraudulent issuers to avoid regulatory restrictions on other forms of fundraising. There are many examples of such frauds, included two cited in Benedetti and Kostovetsky (2018). To our knowledge, the literature has not yet modeled this mechanism formally. This is not out of a blindness to its existence, but because such a model would trivially imply that ICOs should

be banned, without any more nuanced or interesting implications for policy-makers. Moreover, such a model would counterfactually predict a complete breakdown in the market for ICOs (a classic lemons problem), whereas the market seems to be alive and well (and indeed delivering positive *unconditional* returns to investors, based on Benedetti and Kostovetsky, 2018). Conversely, some may even argue it is the credible threat of SEC overregulation that negatively impacted ICO market growth. Without a doubt, the literature should evolve toward delivering clear answers to policymakers on how to deal with fraudulent issuance. A necessary first step is to converge on a convincing model for why an ICO could be attractive *without* fraud. Only then can one fully analyze the trade-offs involved in attempting to screen out the bad ICOs through blunt regulatory means.

4.2 Directions for Future Theoretical Research

Based on the summary above, we identify several potential future directions for theoretical research in the space:

- How to optimally structure ICO deals from a security design perspective. This direction will necessarily involve more careful modeling of the interaction between ICO structuring and project development, and be helpful for linking structural models with reduced-form ones.
- Speculative bubbles: A common concern is that many ICOs represent irrational speculative bubbles, and more theoretical work is needed, including a rigorous definition of “bubble” in the ICO context. Several theoretical papers described above feature fragility and volatility due to multiplicity of equilibria, but this is not the same thing as an irrational bubble.
- Industry classification: given the vast variety of ICO tokens, it would be helpful to provide a framework to classify tokens based on their *economic* functions, rather than *legal* appearances. This direction will closely connect with empirical work in the space.

5 Conclusion

It is rare to witness the birth of a new technology, a new tool for entrepreneurship, or a new investable asset class. Initial coin offerings represent all three, and so it is not surprising that they have sparked much fascination in the finance and technology communities. Undoubtedly, some ICO activity is driven by irrational exuberance. But it is too easy to write off the entire

phenomenon this way. The proper task for academics is to ascertain what kernels of wisdom may lie at the heart of the enthusiasm for this new phenomenon. The research reviewed in this article has made the first steps in addressing this question. We hope that our summary can provide a useful perspective to the next crop of papers that will continue this progress.

Acknowledgements We thank Lin Cong and Sabrina Howell for helpful comments.

References

- Adhami, Saman, Giancarlo Giudici, and Stefano Martinazzi. 2018, forthcoming. Why do businesses go crypto? An empirical analysis of initial coin offerings. *Journal of Economics and Business*.
- Amihud, Y. 2002. Illiquidity and stock returns: Cross-section and time series effects. *Journal of Financial Markets* 5: 31–56.
- Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen. 2006. *Liquidity and asset prices*. Now Publishers Inc.
- Amsden, Ryan, and Denis Schweizer. 2018, June. Are blockchain crowdsales the new ‘gold rush’? Success determinants of initial coin offerings.
- Bakos, Yannis, and Hanna Halaburda. 2018, June. The role of cryptographic tokens and ICOs in fostering platform adoption.
- Benedetti, Hugo, and Leonard Kostovetsky. 2018, May 20. Digital tulips? Returns to investors in initial coin offerings. SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3182169.
- Bourveau, Thomas, Emmanuel T. De George, Atif Ellahie, and Daniele Macciocchi. 2018, July. Initial coin offerings: Early evidence on the role of disclosure in the unregulated crypto market. SSRN: <https://ssrn.com/abstract=3193392>.
- Canidio, Andrea. 2018, August 22. Financial incentives for open source development: The case of Blockchain.
- Catalini, Christian, and Joshua Gans. 2018, June 25. Initial coin offerings and the value of crypto tokens.
- Chod, Jiri, and Evgeny Lyandres. 2018, May 10. A theory of ICOs: Diversification, agency, and information asymmetry.
- Cong, Lin William, and Zhiguo He. 2018, forthcoming. Blockchain disruption and smart contracts. *Review of Financial Studies*.
- Cong, Lin William, Ye Li, and Neng Wang. 2018, July 30. Tokenomics: Dynamic adoption and valuation.
- Deng, Xin, Yen-Teik Lee, and Zhengting Zhong. 2018. Decrypting coin winners: Disclosure quality, governance mechanism, and team networks. SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3247741.

- Fisch, Christian. 2018, forthcoming. Initial coin offerings (ICOs) to finance new ventures. *Journal of Business Venturing*.
- Garratt, Rod, and Maarten van Oordt. 2019, May. Entrepreneurial incentives and the role of initial coin offerings.
- Howell, Sabrina, Marina Niessner, and David Yermack. 2018, June 21. Initial coin offerings: Financing growth with cryptocurrency token sales.
- Hu, Albert, Christine A Parlour, and Uday Rajan. 2018. Cryptocurrencies: Stylized facts on a new investible instrument.
- Kiyotaki, Nobuhiro, and Randall Wright. 1989. On money as a medium of exchange. *Journal of Political Economy* 97 (4): 927–954.
- Lee, Jongsub, Tao Li, and Donghwa Shin. 2018, July. The wisdom of crowds and information cascades in FinTech: Evidence from initial coin offerings.
- Leland, Hayne E., and David H. Pyle. 1977. Informational asymmetries, financial structure, and financial intermediation. *The Journal of Finance* 32 (2): 371–387.
- Li, Jiasun. 2018. Profit sharing: A contracting solution to harness the wisdom of the crowd. Available at SSRN 2844335.
- Li, Jiasun, and Guanxi Yi. 2018, October 1. Toward a factor structure in crypto asset returns. <https://ssrn.com/abstract=3258231> or <https://dx.doi.org/10.2139/ssrn.3258231>.
- Li, Jiasun, and William Mann. 2018. Digital tokens and platform building. Available at SSRN: <https://ssrn.com/abstract=3088726>.
- Momtaz, Paul. 2018, May. Initial coin offerings.
- Olszewski, Wojciech, and Mikhail Safronov. 2018. Efficient cooperation by exchanging favors. *Theoretical Economics* 13 (3): 1191–1231.
- Schelling, Thomas. 1960. *The strategy of conflict*. Cambridge, MA: Harvard University Press.
- Sockin, Michael, and Wei Xiong. 2018, June. A model of cryptocurrencies.
- Townsend, Robert. 1980. Models of money with spatially separated agents. In *Models of monetary economies*, ed. J.H. Kareken and Neil Wallace, 265–304. Minneapolis: Federal Reserve Bank of Minneapolis.
- Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1993. Asset markets as an equilibrium selection mechanism: Coordination failure, game form auctions, and tacit communication. *Games and Economic Behavior* 5 (3): 485–504.
- Yermack, David. 2017. Corporate governance and blockchains. *Review of Finance* 21 (1): 7–31.

Crowdfunding



How Wise are Crowds on Crowdfunding Platforms?

Gilles Chemla and Katrin Tinn

1 Introduction

While firms have obtained outside financing from large numbers of investors on financial markets for centuries, online crowdfunding platforms where crowds of individual, e.g., retail, investors fund projects have only emerged as a major source of funding for start-ups and new projects for a decade. In perhaps the most comprehensive overview of existing platforms to date, Rau (2020) found that funding volumes through about 1300 online platforms rose from around \$0.5 billion in 2011 to about \$140 billion in 2015.

Some platforms such as Kickstarter enable both established firms and, more often, start-ups to pre-sell new products through so-called reward-based crowdfunding campaigns where the only information available to prospective buyers or backers is typically merely a verbal description of the future products and a video showing a prototype. Others give access to a future service whose value often depends on the number of backers using the service. One example is a subset of Ethereum-based contracts in so-called Initial Coin

G. Chemla (✉)

Imperial College Business School, London, UK

e-mail: g.chemla@imperial.ac.uk

K. Tinn

Desautels Faculty of Management, McGill University, Montreal, Canada

e-mail: katrin.tinn@mcgill.ca

Offerings. Security-based platforms, e.g., the Lending Club and AngelList, market more conventional debt-based or equity-based securities but again the information available to investors is typically limited.

This chapter examines the efficiency of investment decisions made by crowds of investors on these various platform types. Whether crowds make efficient decisions is by no means a new question in and outside finance. Outside finance, the argument that crowds can make better decisions than few selected experts dates back to at least Greek philosopher Aristotle's *Politics* in the fourth century BC and it has been widely discussed across centuries and fields (see, e.g., Galton 1907). More recently, the wisdom of crowds was popularized by Surowiecki (2005) through a variety of examples and settings. According to him, for crowds to be wise, they must encompass a diversity of opinion and each person's opinion should be independent of those around them.

In finance, at least since the East India companies in the seventeenth century, financial markets have financed large, risky firms that could not be funded as effectively by a restricted number of investors. Financial markets play an essential role in facilitating coordination by pooling and allocating resources and by generating information about the value of private and public firms. While episodes of misvaluations, bubbles, and crises have attracted a lot of attention (Angeletos et al. 2010; Brunnermeier and Oehmke 2013; Tinn 2010), information generated on marketplaces has broadly proven helpful in valuing investments and guiding financial decision-making (Grossman 1989; Chemla and Hennessy 2014; Goldstein and Yang 2015).

However, there are several reasons why one may expect informational issues on crowdfunding platforms to differ from those in traditional markets. The quality of information available on online platforms is generally much less detailed than that available to financial intermediaries and financial market participants. For example, banks typically have detailed information about the way their clients manage their accounts. Publicly listed firms are required to provide substantial amounts of financial information to investors and they are also scrutinized by different types of professional investors and analysts. While new projects and start-ups of various sizes and risk profiles can be financed on crowdfunding platforms, the limited financial information available to backers who simply need an internet connection, a few dollars, and hardly any financial expertise to scroll user-friendly platforms and back a project may raise doubts on their ability to make wise investment decisions. As individual backers do not have resources or incentives to acquire costly information, this raises the question of to which extent large numbers of backer decisions may provide information about investment quality.

Crowdfunding offers perhaps a unique opportunity to examine whether crowds of backers may make wise decisions and whether those can compare well with those of professional investors. This chapter examines whether crowdfunding harnesses the wisdom of crowds in light of the recent literature. Since whether or not a backer's valuation is independent of other backers' appears to play an important role, it will be useful to organize the discussion around two broad types of crowdfunding: the first one, covered in Sect. 2, where backers' valuation for the item that they buy is private to them and independent of others', as is typically the case when they pre-order a product on platforms such as Kickstarter, and the second one, discussed in Sect. 3, where investor valuations are common, as in the case of long-term investment in a share or a bond, or depend on other investor valuations as in the case when purchasing a service that embodies network effects.

2 Private Valuations in Reward-Based Crowdfunding

On reward-based crowdfunding platforms such as Kickstarter and Indiegogo, firms can raise funds by pre-selling directly future products to potential consumers. Backers then form crowds whenever demand for products offered on the platforms appeals to a large number of potential consumers, as is typically the case for consumer products. In fact, innovative consumer products represent the bulk of projects funded on platforms such as Kickstarter. Chemla and Tinn (2019) found that on August 8, 2018, 65% of the \$3.4 billion raised on that platform funded firms selling technology, design or gaming products; successful campaigns raised \$93 K, \$63 K and \$55 K in technology, design and games on average, respectively, but only \$10 K across all other categories; 284 out of the 310 projects that had raised over \$1 million (4,147 out of the 5,265 projects that had raised over \$100 K) were in the three categories mentioned above.

One may wonder why backers will pre-order new products rather than invest directly in the project by buying securities. Earlier explanations have relied on backer preferences. One argument is that participating in the development of a new product confers backers an increase in their utility relative to their valuation for the product itself (Belleflamme et al. 2014; Varian 2013). This argument is compelling and it is certainly valid for some types of projects, e.g., with innovative, fun, or social characteristics.

This argument is unlikely to apply for most projects, however. If backers obtain an additional utility from funding a project, they will be willing to

pay more for the product during a crowdfunding campaign than the future consumers who will buy it after the product development phase has been completed. This will give firms an incentive price-discriminate and to choose a crowdfunding price that is higher than the future retail price. However, evidence from Kickstarter suggests that crowdfunding prices are often lower than retail prices later on. In addition, the argument developed in Belleflamme et al. and Varian implies that backers may not be representative of the population of consumers and hence that corporate investment decisions based on backer strategies may be biased. Of course, price discrimination incentives may also emerge in other contexts, e.g., when firms do not know interested backers' willingness to pay (see Ellmann and Hurkens 2019), a consideration that may also affect the empirical assessment of crowdfunding outcomes.

Absent considerations of an additional backer utility derived from participating in crowdfunding, that each backer's valuation for the products pre-sold on reward-based crowdfunding platforms is private and independent of other backer valuations can be very valuable for firms. In such an environment, Chemla and Tinn (2019) point out that reward-based crowdfunding platforms enable firms to obtain a proof of concept at an early stage of product development. The argument goes as follows: Pre-selling a product on platforms acts as a credible consumer survey where firms learn about target consumer preferences and demand before making investment decisions. This creates a substantial real option value: observing the decisions of a representative sub-sample of consumers enables firms to obtain information about demand from all future consumers.

In fact, both success and failure provide helpful information to firms: Firms may benefit from learning that their product will enjoy high consumer demand, which will prompt them to invest, or that the product will not sell in large enough numbers, in which case information enables firms to save on investment costs. The real option value of learning is particularly high when firms are more uncertain about future demand and when their investment cost is close to the present value of cash flows based on prior beliefs, in which case information about demand is particularly valuable. Innovative consumer products with high demand uncertainty are likely to benefit most from reward-based crowdfunding. From that point of view, backer crowds will help firms to obtain information that will lead them to make wise decisions.

Empirical evidence is very much consistent with these findings. Viotto da Cruz (2018) and Xu (2017), among others, find that entrepreneurs are more

likely to complete their projects after the level of pledging through crowdfunding increases and that the effects are stronger when projects are subject to high demand uncertainty and when the crowd is more experienced. Xu also finds that entrepreneurs tend to start riskier projects when the opportunity cost of crowdfunding increases, which is consistent with the real option value of crowdfunding.

However, crowdfunding may be subject to a well-known moral hazard problem endemic in corporate finance (Tirole 2006) and previously discussed in Strausz (2017) in reward-based crowdfunding. Despite the real option value of learning, firms may be tempted to divert the funds raised during a campaign instead of completing a project. This problem may be all the more acute as reward-based crowdfunding platforms are generally not legally responsible for the delivery of products. Further, providing compelling evidence that a firm has committed a fraud is difficult and coordinating a legal action among backers may be costly and complicated. Hence, if the funds collected through crowdfunding and the number of backers are large, firms may be more tempted to divert funds.

Yet, Mollick (2014) finds that only 3.6% of Kickstarter projects have failed to deliver products following a successful campaign. Further, the projects that receive noticeably more funds than their targets do not appear to fail to deliver rewards more frequently than others. This may seem surprising at first. However, Chemla and Tinn (2019) show that the real option value of learning enables firms to overcome moral hazard and reduces the incentives for firms that raise more funds typically from a large number of backers to divert them. The reason is that after its crowdfunding campaign, a firm that expects there to be sufficiently high future consumer demand will choose to not divert funds even if it is costless to do so. During a crowdfunding campaign, firms will often choose a target of funds to raise that is set so as to make sure that corporate investment is incentive compatible. Specifically, campaigns will be designed so that the platform will return the funds to backers if that target is not met so firms will have no funds to divert. This so-called “All-or-Nothing” (AoN) crowdfunding scheme with a sufficiently high target makes sure that firms have incentives to invest after a successful campaign.

One may then wonder whether large enough crowds and, relatedly, long enough crowdfunding campaigns are desirable. It turns out that because of the threat of moral hazard, shorter campaigns are more likely to succeed and a limited backer sample maximizes both expected fundraising and platform fees. The reason for this is simple. If the product was entirely pre-sold during the campaign, then firms would always be tempted to divert funds instead of

delivering the products. In fact, what makes the completion of the project incentive compatible is the possibility for firms to sell to a large enough number of consumers after both the campaign and the investment decision. Overall, the crowdfunding sample or the crowd must be sufficiently small. However, since firms crowdfunding their own project on their own platform may be tempted to extend the campaign length, third-party platforms may be preferable to make sure that the crowdfunding sample is indeed credibly small enough.

One may then wonder how crowdfunding compares with an investment in an entrepreneurial project that enables a firm to experiment a new technology (Hellmann 2002; Bettignies and Chemla 2008; Kerr et al. 2014). It turns out that crowdfunding has the advantage of enabling firms to learn about demand without having to bear investment costs.

Crowdfunding comes with another important advantage: the information that it generates makes it complementary to other, more traditional forms of outside financing. In several cases, after successful reward-based crowdfunding campaigns, firms obtain further funding from angels, venture capitalists, and investor-based crowdfunding which they had not obtained before the crowdfunding campaign. In such cases, the information generated by a crowd of backers may be helpful to a small group of professional investors.

3 Common Values and Securities-Based Crowdfunding

Although the pre-sale of products in private value settings is an important feature of the projects crowdfunded on prominent platforms such as Kickstarter and Indiegogo, crowdfunding platforms marketing securities and projects with imperfect information about product quality or network effects lead to backer valuations that involve common values. In such environments, each backer's valuation for a project being crowdfunded depends on their beliefs about other backer valuations. This may lead backers to pay excessive attention to other backers' strategies and generate coordination problems and informational cascades.

For example, Li and Mann (2018) explore network effects in the context of the subset of Initial Coin Offerings (ICOs) that resemble reward-based crowdfunding in that the investors offer funds in exchange for access to a future service that the firm is developing. If the private value of the service increases with the number of backers who will use that service, pre-selling

can help coordination. Many other forms of crowdfunding and ICOs involve investors obtaining financial rewards that are common across investors.

When considering securities, larger crowds can still sometimes make superior predictions. Iyer et al. (2015) find some evidence consistent with the wisdom of crowds in online lending platforms: Peer lenders who cannot observe credit scores can still predict an individual borrower's likelihood of defaulting with 45% greater accuracy than the borrower's credit score. Even when they cannot beat experts, peer lenders can achieve astonishing results. For example, they achieve 87% of the predictive power of an econometrician who observes all standard financial information about borrowers. This suggests again some complementarity between platforms and traditional financiers as combining information from both sources can enhance efficiency on the lending market.

Yet, the evidence regarding lending platforms is mixed and does not appear to always harness the wisdom of crowds. Mohammadi (2019) compares the crowd's ability to screen the creditworthiness of small-and medium-sized enterprises (SMEs) who applied for loans on the FundingCircle peer-to-peer lending platform with that of institutions. He finds that crowd underperforms institutions in screening SMEs and that they fail to lend at interest rates that adjust for the likelihood of defaulting. The interest rate set by crowd predicts default 39% less accurately than institutions. Further, there is also some evidence that lending platforms allow sophisticated investors to take advantage of unsophisticated investors, Hildebrand et al. (2016) find that group leader bids in the presence of origination fees are wrongly perceived by other lenders as a signal of good loan quality.

The sensitivity of a backer's belief about a project to others' valuations plays an important role in the debate on the wisdom of crowds and informational cascades. Just like in financial markets, this sensitivity is arguably more pronounced for information-sensitive securities such as equity. Brown and Davis (2019) examine a securities-based crowdfunding environment where dispersed, privately informed investors simultaneously decide whether or not to back a project. They find that an optimal offering cannot generate a wisdom of the crowd result. Indeed, the investment decision by the entrepreneur affects investors' incentives to invest truthfully based on their private information. Cong and Xiao (2018) examine the informational efficiency of securities-based crowdfunding when investors move sequentially and they show that all-or-nothing thresholds improve the informational efficiency of crowdfunding.

Astebro et al. (2019) study the micro-level data of herding on Seedrs, a major European equity crowdfunding platform. They find that when the size

of a pledge made by a backer doubles, a subsequent pledge is between 5.8 and 20% larger. This is consistent with the idea that a large pledge signals that the backer making the pledge has received positive information about the project that others might not. This in turn may cause follow-on investors to alter their investment strategies, even though they don't actually observe the information available to the investor making the large pledge. They also find that the probability that a campaign is successful depends largely on the support it gets early in the campaign. This may be due to the effect that low pledges have on investor beliefs about the project. A lack of support for a campaign is indicative that only a few investors are arriving with positive signals. Hence, having little support early in the campaign makes potential backers more pessimistic that the project is of good quality. Investors then either pledge lower amounts or decide to not invest. In this context an abstention information cascade is likely to occur, and failed campaigns end up missing the mark by a large margin. Conversely, broad support early in the campaign may lead to excessive campaign success.

4 Concluding Remarks

Private value environments such as those on reward-based crowdfunding platforms such as Kickstarter exhibit patterns that are consistent with the wisdom of crowds. Such platforms enable firms to experiment and learn valuable information about demand which creates a valuable real option. However, it may be beneficial to limit the size of crowds on such platforms so meeting future consumer demand be an important concern to firms and prompt them to complete their project and produce and deliver the products pre-ordered by backers.

In many other environments, backer valuations depend on other backer strategies. This is the case, for example, when ICOs monetize a provision of services with network effects and when backers buy securities. Some evidence suggests that lending platforms may predict default probabilities better than common credit scoring instruments. However, beliefs about other backer beliefs will often entail fragilities prone to generate informational cascades. Further, investor bias such as home bias (Lin and Viswanathan 2019) and the performance of textual analysis and various types of algorithms (Dorfleitner et al. 2016; Jiang et al. 2018; Jain and Jamieson 2018) that are arguably more easily accessible to professionals than to individual investors may also limit the ability of crowdfunding platforms to harness the wisdom of crowds. This will certainly be the subject of much research.

References

- Angeletos, G., G. Lorenzoni, and A. Pavan. 2010. *Beauty contests and irrational exuberance: A neoclassical approach*, MIT.
- Astebro, T. B., Fernández Sierra, M., Lovo, S., and Vulkan, N. (2019, June). Herding in equity crowdfunding. In Paris December 2018 Finance Meeting EUROFIDAI-AFFI.
- Belleflamme, P., T. Lambert, and A. Schwienbacher. 2014. Crowdfunding: Tapping the right crowd. *Journal of Business Venturing* 29: 585–609.
- Bettignies, J., and G. Chemla. 2008. Corporate venture capital, allocation of talent, and competition for star managers. *Management Science* 54: 505–521.
- Brown, D.C., and S. Davis. 2019. Financing efficiency of securities-based crowdfunding. *Review of Financial Studies*, forthcoming.
- Brunnermeier, M.K., and M. Oehmke. 2013. Bubbles, financial crises, and systemic risk. In *Handbook of the Economics of Finance* 2: 1221–1288. Elsevier.
- Chemla, G., and C.A. Hennessy. 2014. Skin in the game and moral hazard. *The Journal of Finance* 69 (4): 1597–1641.
- Chemla, G., and K. Tinn. 2019. Learning through crowdfunding. *Management Science*, in press.
- Cong, L., and Y. Xiao. 2018. *Information cascades and threshold implementation*. Available at SSRN: <https://ssrn.com/abstract=3030573>.
- Dorflleitner, G., C. Priberny, S. Schuster, J. Stoiber, M. Weber, I. de Castro, and J. Kammler. 2016. Description-text related soft information in peer-to-peer lending—evidence from two leading European platforms. *Journal of Banking & Finance* 64: 169–187.
- Ellman, R., and S. Hurkens. 2019. Optimal crowdfunding design. *Journal of Economic Theory*, in press.
- Galton, F. 1907. Vox populi (the wisdom of crowds). *Nature* 75 (7): 450–451.
- Goldstein, Itay, and Liyan Yang. 2015. Information diversity and complementarities in trading and information acquisition. *The Journal of Finance* 70 (4): 1723–1765.
- Grossman, S. 1989. *The informational role of prices*. MIT Press Books.
- Hellmann, T. 2002. A theory of strategic venture investing. *Journal of Financial Economics* 64: 285–314.
- Hildebrand, T., M. Puri, and J. Rocholl. 2016. Averse incentives in crowdfunding. *Management Science* 63 (3): 587–608.
- Iyer, R., A. Khwaja, E. Luttmer, and K. Shue. 2015. Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62 (6): 1554–1577.
- Jain, L., and K. Jamieson. 2018. Firing bandits: Optimizing crowdfunding. In *International Conference on Machine Learning*, 2211–2219.
- Jiang, C., Z. Wang, R. Wang, and Y. Ding. 2018. Loan default prediction by combining soft information extracted from descriptive text in online peer-to-peer lending. *Annals of Operations Research* 266 (1–2): 511–529.

- Kerr, W., R. Nanda, and M. Rhodes-Kropf. 2014. Entrepreneurship as experimentation. *Journal of Economic Perspectives* 28 (3): 25–48.
- Li, J., and W. Mann. 2018. *Initial coin offerings and platform building*. Available at SSRN: <https://ssrn.com/abstract=3088726>.
- Lin, M., and S. Viswanathan. 2019. Home bias in online investments: An empirical study of an online crowdfunding market. *Management Science*, forthcoming.
- Mohammadi, A. 2019. *How wise are crowd? A comparative study of crowd and institutions in peer-to-business online lending markets*. Mimeo University of Gottenburg.
- Mollick, E. 2014. The dynamics of crowdfunding: An exploratory study. *Journal of Private Equity* 29: 1–16.
- Rau, P. Raghavendra. (2020, July). “Law, trust, and the development of crowdfunding.” Trust, and the Development of Crowdfunding.
- Surowiecki, J. 2005. *The wisdom of crowds*. Anchor.
- Strausz, R. 2017. A theory of crowdfunding—A mechanism design approach with demand uncertainty and moral hazard. *American Economic Review* 107 (6): 1430–1476.
- Tinn, K. 2010. Technology adoption with exit in imperfectly informed equity markets. *American Economic Review* 100 (3): 925–957.
- Tirole, J. 2006. *The theory of corporate finance*. Princeton University Press.
- Varian, H. 2013. *Public goods and private gifts*. Mimeo UC Berkeley.
- Viotto Da Cruz, J. 2018. Beyond financing: Crowdfunding as an informational mechanism. *Journal of Business Venturing* 33 (3): 371–393.
- Vismara, S. 2016. Information cascades among investors in equity crowdfunding. *Entrepreneurship theory and practice*.
- Xu, T. 2017. *Learning from the crowd: The feedback value of crowdfunding*. Mimeo UBC.



Asymmetric Information at the Platform Level: The Case of Peer-To-Peer Lending

Puriya Abbassi and Raj Iyer

1 Asymmetric Information in Credit Markets

Channelling funds in financial markets is widely known to involve two main issues: one is related to search costs and the matching of borrowers with lenders and vice versa; another issue, which is probably more severe, is associated with the presence of asymmetric information, where one party does not know enough about the other party to make accurate decisions. The asymmetry in information creates problems in credit markets on two fronts: before the transaction is entered into and after.

Before the transaction is made, a key friction occurs when potential borrowers who are the most likely to produce an undesirable (adverse) outcome, i.e., the bad credit risks, are the ones who most actively seek out a loan and are thus most likely to be selected (adverse selection). Because adverse selection makes it more likely that loans might be made to bad credit

The views expressed in this chapter are solely those of the author(s) and do not necessarily reflect those of the Deutsche Bundesbank or the Eurosystem.

P. Abbassi

Deutsche Bundesbank, Frankfurt am Main, Germany

e-mail: puriya.abbassi@bundesbank.de

R. Iyer (✉)

Imperial College London, London, UK

e-mail: r.iyer@imperial.ac.uk

risks, lenders may decide not to make any loans even though there are good credit risks in the marketplace, a key concept advocated by Akerlof (1970). Another friction arises after the transaction is made, i.e., when the lender faces the risk that the borrower might engage in activities that are undesirable (moral hazard) from the lender's perspective as they make it less likely that the loan will be paid back. Since moral hazard lowers the probability that the loan will be repaid, lenders may decide not to make a loan. As a result, even projects with a positive net present value might not get funded. Moreover, as Stiglitz and Weiss (1981) point out, imperfect information in lending markets can leave some borrowers unfunded even if they are willing to pay a price higher than the market interest rate.

There is a rich body of theoretical and empirical research that focuses on how these problems associated with imperfect information can be mitigated. Gorton and Winton (2003) review empirical evidence for why financial intermediaries exist, and suggest several roles that banks exhibit. Two of these roles are directly related to the issue of information asymmetry—banks as information producers and banks as delegated monitors. Diamond (1984) argues that banks exist due to the high costs of ex-post monitoring of borrowers; hence, intermediaries are established to economize on these activities. This in turn raises the question as to who 'monitors the monitors', where lenders need to be assured that the intermediary itself is indeed monitoring the borrowers. Williamson (1987) discusses the important role of diversification in alleviating this problem. If the intermediary has an infinitely large number of borrowers, the lenders will not need to monitor the intermediary as the market return can be ensured with certainty (Gorton and Winton 2003).

Another important reason for the existence of financial intermediaries, as argued by Gorton and Winton (2003), is that without intermediaries, lenders would be faced with two inherent issues of information production. One refers to the fact that a small number of agents can produce information about the borrowers and sell it to other lenders, thereby economizing the cost of information production. However, these entities cannot credibly ensure that the information produced is valuable (Gorton and Winton 2003). That is, even if the information producer can assure other agents of the information's value, purchasers of the information can easily re-sell such information, or share it with others. This would further render the production of information less economic (Hirshleifer 1971; Grossman and Stiglitz 1980; Gorton and Winton 2003).

A key insight from all of these theoretical and empirical contributions is that financial intermediaries are appropriately designed and equipped to screen risks ex-ante and to monitor borrower's (loan) performance ex-post,

thereby mitigating problems of information asymmetries and facilitating a more efficient allocation of capital (Diamond 1984; Blackwell and Winters 1997; Boot 1999; Boot and Thakor 2000).

Another key insight is that banks do very little screening for small borrowers, but instead rely heavily on collateral. Providing collateral, however, is costly and binds capital such that some creditworthy borrowers with higher shadow cost of capital are prevented from obtaining loans (Stiglitz and Weiss 1981; Ang et al. 1995; Avery 1998; Manove et al. 2001). In fact, Berger et al. (2005) show that large banks are less willing to lend to firms with no financial records. Moreover, after controlling for the endogeneity of bank-firm matching, they find that large banks lend at a greater distance, interact more impersonally with their borrowers, have shorter and less exclusive relationships, and do not alleviate credit constraints as effectively. This suggests hampered access to credit markets for non-incumbent borrowers. While micro credit institutions have improved financial access for small borrowers and borrowers with missing financial records in many economies, they primarily rely on group lending principles, which can make it difficult for individual borrowers to access credit (Iyer et al. 2016). And as equity and corporate debt markets also typically provide financing to large-scale, mature borrowers, direct, market-based finance cannot substitute the impeded access to credit for smaller borrowers in the traditional bank lending market.

2 The Rise of Online Credit Markets and Peer-To-Peer Lending

The growth in digitization and Internet technologies has substantially altered the way online network users interact and connect themselves with each other (Lin 2009). Increasingly, these online networks are allowing for new business models, thereby creating an environment in which traditional commerce in several industries undergoes a transformative process. One of these industries that has undergone transformation is the market for financial lending, in which a new version of the traditional private lending market has re-emerged. These online credit markets enable private lenders and borrowers to engage with each other and to undertake loan origination without the mediating involvement of financial institutions. Therefore, these online credit markets are often referred to as peer-to-peer lending markets. Thus, the online model of private lending requires a mere platform which allows for the interaction

of private borrowers and lenders to facilitating search and matching. In addition to their loan request, borrowers can use the platform to provide an array of additional information, including the project for which they wish to seek financing, their current financial situation (e.g., income, credit score, etc.) as well as personal details (e.g., age, gender, etc.). Based on the provided information, lenders initially could select a borrower, derive an interest rate and make an offer. Nowadays, however, platforms require lenders to auto-select loans, i.e., lenders no longer can choose to which borrower to lend.

The first lending platform that created this kind of forum for private lenders and borrowers was 'Zopa' (UK), established in 2005. Ever since, a variety of platforms emerged with about 52 distinct business models existing worldwide by the end of 2018.¹ The Cambridge Center for Alternative Finance (CCAF) identifies over 3,500 platforms in 182 countries (Rau 2019). All of the platforms generate their revenue by imposing certain fees on the intermediation of the transaction. These online lending markets have grown in popularity and size.² These platforms cover a wide range of market segments, including consumer credit, loans to small businesses, student loans and real estate. The uncollateralized nature of lending in these online credit markets makes it particularly attractive for small borrowers who otherwise might turn to payday lenders or credit card debt, often at exorbitant rates (Adams et al. 2009; Iyer et al. 2016).

Initially, platforms with a commercial focus operated nationally, while non-commercial online lending platforms served borrowers and lenders globally (Ashta and Assadi 2009). Nowadays, however, lend-a-hand and other platforms are increasingly operating both commercially and internationally. There are also different ways upon which platforms conduct the loan origination process. For instance, while some connect lenders and borrowers directly, others bring them together via a third party, usually a bank (Bachmann et al. 2011). Furthermore, platforms differ in the way the loan's interest rate is set. Platforms like 'Prosper' (US) use an auction-type setting, whereas others such as 'Smave' (Germany) compute the rate based on borrower characteristics, both financial and demographic. Most of the platforms also provide features that facilitate the use of social networks. More precisely, they encourage borrowers and lenders to form online groups, and establish friendships with other members. Leaders and friends can offer endorsements for a specific listing and highlight bids from group members, endorsing group leaders, and friends.

¹ Rau (2019) describes the major types of business models reported by platforms on a global basis.

² Refer to www.p2p-banking.com for further information.

3 Information Asymmetries in Peer-To-Peer Lending

As non-financial experts dominate these online lending markets, their ability to judge financial risk and information becomes crucial for the sustainability of these markets, and thus their effectiveness to allocate credit to smaller and less mature borrowers. In fact, since many of these peer-to-peer platforms keep individual users anonymous to each other, it seems reasonable that the extent of information asymmetry is likely to be greater on online credit markets than on offline traditional credit markets. Therefore, a crucial question is as to how online lending markets cope with frictions associated with imperfect information.

Generally, peer-to-peer lending market listings require three types of borrower information. First, some online lending platforms cooperate with credit bureaus and collection agencies, and therefore draw on borrower's credit history. Second, the peer-to-peer lending websites require any borrower seeking to use the platform to provide certain social network variables. And third, some platforms allow borrowers to upload additional self-reported non-verifiable information.

To screen a borrower's risk, the traditional loan officer relies strongly on borrower's actual credit score ('hard information'), especially for small borrowers. Lenders operating on peer-to-peer lending websites, however, can only observe—if at all—credit *grades*. The difference being, scores are continuous, and thus finer with respect to the associated credit risk assessment. Grades, however, are discrete, and thus as fine as the credit grade width allows. While this information allows for discriminatory pricing of loans to borrowers across credit grades, Freedman and Jin (2011 and 2017) show that categorical credit assessment creates adverse selection within each grade. Moreover, Iyer et al. (2016) expand the analysis to an array of 'hard information' posted by the borrower, such as debt-to-income ratio, borrower's current number of delinquencies, amount delinquent, and the number of credit inquiries in the last six months, and show that lenders in such peer-to-peer markets are able to use borrower information to infer a third of the variation in creditworthiness within a credit category that is captured by a borrower's credit score. However, they also find that lenders learn about borrower's creditworthiness by using non-standard subjective information obtained from borrowers' self-posted profiles.

Freedman and Jin (2011 and 2017) examine the information content of 'soft information' requested by peer-to-peer lending websites, using the example of 'Prosper'. More precisely, they study how information conveyed

from social networks may serve as signalling device, and hence help lenders to screen borrower's risk before engaging in a lending relationship (Spence 1973). For instance, friends and group leaders may certify a borrower's history of on-time payment, and some social connections reflecting education and employment may be good indicators of borrower's payment ability. Likewise, the fear of reputational damage in front of friends and group members may act as an additional motivation for the borrower to repay the loan. The information content is not necessarily straightforward as some social networks may operate differently. For instance, a group lender may endorse a risky borrower because she seeks to fetch the group leader award and therefore does not engage in a serious and adequate risk screening. Moreover, a borrower may strategically ask a friend to endorse and bid on her listing against some small financial compensation. Freedman and Jin (2011 and 2017) use transaction-level data on loans originated on 'Prosper' for the period 2006:M6 to 2008:M7 and show that soft information conveyed from several social networks can help to mitigate information asymmetry while others do not, depending on the institutional incentives of the respective network. For example, loans that had both friends' endorsement and friend bids are associated with fewer missed payments than others. However, loans, where friends endorsed the inquiry but no friend bid, are associated with a lower rate of return than loans without endorsement. On average though, default rates of group loans are higher than non-group loans. And, while group ratings add little information, there is evidence on some group lenders to strategically extract group awards without adequately screening the borrower's risk. In fact, Hildebrand et al. (2016) show that group leader bids in the presence of origination fees are ex-post associated with higher default rates suggesting that these bids are wrongly perceived as a signal of good loan quality.

Lin et al. (2013) also study the relationship between online friendship and transactional outcome on peer-to-peer lending platforms. They find that borrowers with friends are more likely to have their loan requests funded and that these loans have lower interest rates. Moreover, they find that friendships lower default probability. But both the ex-ante and ex-post results show a striking gradation along friend type, with greater effects when friends have roles and identities that signal better credit quality. This suggests that friendships serve as informational cues of a borrower's credit quality. More generally, the authors support the central premise of signalling models that agents facing asymmetric information adapt by using signals to mitigate adverse selection (Spence 1973). The results also highlight how technology aids this process by facilitating the generation and transmission of new sources of information.

In the literature, disclosure of non-verifiable information has been discussed to affect peer-to-peer lending (e.g., Michels 2012). Ravina (2018) extends the analysis and examines whether and how borrower's personal characteristics, i.e., borrower's self-reported non-verifiable information, affect lenders decision to grant a loan. The intuition of the study being, lenders make inferences based on past experience, and therefore base their credit risk assessment on easily observable variables that have already proven to be correlated with ex-post performance in the past. These credit risk assessments involve statistical discrimination models, where the group that has been categorized as less creditworthy is less likely to get a loan, pays a higher interest rate and is associated with higher likelihood to underperform as compared to a group that is considered more creditworthy.³ However, one may argue that these easily observable characteristics do not have additional explanatory power to predict ex-post performance once reliable and verifiable information has been taken into account (Becker 1957). In such a case, lenders may be even willing to accept losses in profits just in order for them not to interact with certain groups that they deem reasonable to discriminate. Similarly, lenders may discriminate groups depending on their subjective perception of the borrower's health, intelligence, and beauty, thereby making inference about their competence to pay back the loan. Other studies, however, point out that identity claims constructed in narratives by borrowers have also a significant impact on lender decisions about unsecured personal loans (e.g., Herzenstein et al. 2011). In fact, Herzenstein et al. (2011) show that some non-verifiable information are provided to intentionally to mislead lenders.

Using transaction-level loan data obtained from 'Prosper' over the period 2007:M3 to 2007:M6, Ravina (2018) shows a differential treatment of borrower by lenders, depending on the borrower's gender, race and level of perceived beauty. That is, borrowers that are not good-looking are less likely to receive a loan, and tend to pay higher interest rates, even though they are associated with lower default frequencies than alleged 'ugly' borrowers. Also, she finds that female applicants are more likely to get a loan than male applicants, all else equal. Moreover, she finds that lenders belonging to areas with high racial prejudice against blacks discriminate black borrowers by imposing higher interest rates, even though these black borrowers are less likely to default than black borrowers generally. Furthermore, she shows that black borrowers with a higher fraction of black lenders are more likely to receive a loan, pay less and are less likely to default. That is, she shows that lenders perform better at screening borrowers that belong to their same group, both

³ Such models are widely used in the literature on labor economics, see e.g., Phelps (1972), Arrow (1973), Calomiris et al. (1994) and Lundberg and Startz (1998) (among others).

at the extensive margin (granting the loan) and the intensive margin (lower rates, lower default rates and higher return). These results are consistent with the findings of Pope and Snyder (2011) and Duarte et al. (2012). Pope and Snyder (2011), for instance, also use transaction data from 'Prosper' and find that loan listings with blacks in the attached picture are 25–33% less likely to get the loan than those of whites with similar credit profiles. On the intensive margin, they show that blacks are charged higher average interest rates. Lenders making such loans, however, are found to earn a lower net return compared to loans made to whites with similar credit profiles because blacks have higher relative default rates. Duarte et al. (2012) find that borrowers who appear more trustworthy have higher probabilities of having their loans funded. Moreover, borrowers who appear more trustworthy indeed have better credit scores and default less often.

4 Differences Across Peer-To-Peer Lending Markets

We have established that different platforms provide different environments for borrowers and lenders to engage with each other. A crucial question that arises in this context is how different platforms function compared to each other, from a lender's and a borrower's perspective, respectively. To that aim, we compare the two widely known and most established peer-to-peer lending websites, i.e., 'Prosper' and 'LendingClub'.

While both platforms rely heavily on hard information to assess a borrower's credit grade, using especially income and credit history information, they use different levels of granularity to construct credit grades. For instance, while 'Prosper' uses seven distinct grades of borrowers, 'LendingClub' relies on five credit categories. Another difference among online credit markets refers to the fees they impose while processing a loan origination. Both 'LendingClub' and 'Prosper' typically charge borrowers an origination fee of 5% per annum after a loan's approval. Yet, 'LendingClub' incentivizes more creditworthy borrowers by lowering the origination fee to 1% per annum if they hold an excellent credit history. Similarly high-rated borrowers on 'Prosper', however, face a lowered fee of 2.4% per annum. In contrast to 'Prosper', 'Lending Club' allows borrowers to pool their loan request with a friend or co-applicant, which may boost the overall credit grade, thereby increasing the likelihood to get a loan in the extensive margin, and lower rates in the intensive margin.

For lenders, there are also different options online lending markets offer. For instance, the ease with which an investor can navigate on each website might play a role. Another aspect is the heterogeneity across potential borrowers and the size of the platform's pool of loan applicants, which will have implications for risk-return calculations. Also, some platforms allow loan units only as a function of a minimum loan amount. For instance, while 'LendingClub' allows investments as multiples of \$25, 'Prosper' grants more flexibility with respect to the loan investments as it only requires a loan to be at least \$25.

5 Peer-To-Peer Lending Markets vs. Traditional Lending Markets

While there is an increasing growth in popularity of funding via online lending markets, an important question refers to the type of borrower that these platforms seem to serve. Addressing this question is important for understanding and evaluating the expansion in online lending markets. There is an increasing body of research that seeks to address this question, yet providing mixed evidence on the type of borrowers served by P2P platforms. In particular, these studies examine the different borrower segments operating on non-bank lending markets as compared to the pool of borrowers served by banks, i.e., the traditional bank lending market.

For instance, Buchak et al. (2018) compare US residential mortgage loan originations by shadow banks with those by traditional banks and document a segmented borrower pool, where FinTech lenders seem to attract a larger share of less creditworthy borrowers than banks do. In contrast, Fuster et al. (2018) find no evidence that FinTech lenders target risky or marginal borrowers, and thus suggest that technological innovation has improved the efficiency of financial intermediation in the US mortgage market.

For borrowers in the consumer lending segment, De Roure et al. (2018) show that peer-to-peer loans in Germany are on average riskier than those granted by banks. Nevertheless, once having adjusted the loan rates by the borrower-inherent risk, they find that interest rates on peer-to-peer loans are lower than those reported for bank loans to borrowers of similar risk. Using survey data, Liao et al. (2017) similarly find that, in China, online lending platforms focus on unjustifiable borrowers. In the same vein, contrary findings have been documented in the US consumer credit market by Wolfe and Yoo (2017). They document rather competitive effects suggesting that online lending markets may act as substitutes for borrowers who otherwise

would draw on bank lending. More precisely, they show that small (rural) commercial banks lose lending volume to peer-to-peer lending websites.

Tang (2018) exploits a negative shock to bank credit supply in the United States and shows that peer-to-peer consumer lending expands in the market affected by the negative shock. Moreover, she finds evidence for substitutional effects between banks and online credit markets, notably 'LendingClub', insofar as the quality pool of borrowers in online credit markets deteriorates when low-quality bank borrowers migrate to peer-to-peer lending platforms. This result suggests that the credit expansion opportunities associated with peer-to-peer lenders only benefit infra-marginal bank borrowers. At the same time, however, the peer-to-peer platform 'LendingClub' seems to complement banks by providing small loans. Based on data from 'Prosper', Balyuk (2018) provides evidence that supports the notion that peer-to-peer lending may even relieve information frictions in the consumer credit market and increases access to credit for borrowers even from existing lenders.

While these studies help to improve our understanding of the borrower pool attracted by online credit platforms, it still remains an avenue for future research to explore more comprehensively the borrower segments across different online market segments depending on the loan type, the investor base and the geographical decomposition.

6 Challenges Ahead

Despite growing popularity of peer-to-peer lending platforms worldwide, there are two broad sets of challenges, which need rigorous attention and consideration. One of these challenges relates to the existence and progress of peer-to-peer markets in general. In essence, it relates to the question as to why peer-to-peer markets are experiencing such a growth in recent years to begin with. The answer to this question will have important implications for how peer-to-peer lenders may have to reorganize themselves to play their role in future financial markets.

For instance, one reason for the growing importance of peer-to-peer markets is often linked to banks' reduced lending appetite vis-a-vis small businesses in response to the financial crises 2007/2009. Additionally, the ultra-low interest rate environment induced by central banks worldwide is also considered to add appeal to peer-to-peer investments, especially when traditional savings vehicles of any kind deliver only minimal yields. This, however, suggests that emerging providers of peer-to-peer finance will face

a strong and competitive environment if banks regain their interest in small business financing and/or yields start hiking again.

Furthermore, the recent experience with one of the UK's biggest peer-to-peer lenders unravels more challenges for the FinTech credit sector. In 2014, RateSetter engaged in wholesale lending with another lending company, Vehicle Trading Group. Vehicle Trading Group then lent out the funds further to an advertising firm called AdPod (a client of another peer-to-peer lender). The information that investor's funds at RateSetter were 'passed' on to AdPod became only available in early summer of 2017, i.e., when Vehicle Trading Group went bankrupt. Even though RateSetter was able to take over AdPod and thereby prevent investors from losing their invested funds, there are three important insights one can draw from this experience. First, implementation of rules on plans for wind-downs of peer-to-peer platforms will be important to establish a minimum amount of trust in this sector. Second, investors should be provided with sensible information about how their funds are allocated. Third, lending platforms need a framework that is resilient during shocks and less prone to funding freezes. These claims are, in fact, also raised by peer-to-peer lenders themselves (e.g., P2PFA 2016).⁴

Another important issue with peer-to-peer markets is their exposure to borrowers mostly affected by swings in credit risk in financial markets. Since most of the peer-to-peer clients are either borrowers with low or no credit history, it is reasonable to assume that borrowers from these markets are the first to fail to meet the terms of their lending deals should markets move adversely. Needless to say, a higher proportion of non-performing loans will directly affect the peer-to-peer lender's returns. Zopa, another big UK peer-to-peer lender, for instance, has recently warned its investors that they may see their returns cut for products with higher projected rates of interest due to a rise in consumer debts going bad. While prudent from a risk management perspective, the strategy to focus more on low-risk borrowers adds two more challenges to the FinTech credit business: competition (with banks) and the challenge to acquire new customers. But to meaningfully challenge banks' market position in lending, peer-to-peer lenders may have to become more efficient and more targeted in terms of how they intermediate between borrowers and lenders and thus put pressure on incumbent banks.

Certainly, this is a great opportunity for the financial system to become more diverse and competitive, thereby offering customers better-targeted,

⁴ In the UK, the Financial Conduct Authority (FCA) is making progress on these grounds; new rules reflecting some of these issues will become effective as of December 9, 2019, see <https://www.fca.org.uk/news/press-releases/fca-confirms-new-rules-p2p-platforms>. However, for effective and efficient regulation, international collaboration will be required as we will argue further below.

faster and cheaper financial services. This, however, requires lending platforms to perform liquidity transformation on a significant scale. Two important observations support the notion that the FinTech credit landscape has already started this transition. First, an increasing number of globally active technology companies (big tech companies) are entering the market (e.g., Ant Financial (Alibaba) and Webank, among many others), some of which may change current market segments entirely (e.g., Facebook's recent attempt to introduce its own currency 'libra'). Second, smaller peer-to-peer lenders are starting collaborations with big tech companies to leverage up their community engagement. Recent trends in Kenya and South Africa are important examples of this kind. Having recently partnered with the local value-added reseller Khonology, White Label Crowdfunding (WLCF), is looking to collaborate with the founders of new African platforms and is keen to support the shaping of the market.

These trends, however, raise the second set of new challenges, notably for legislative and regulatory bodies. First and foremost, it is important to understand that these big tech companies are generally less motivated by the return on their credit activity and more by the access they get to data that let them sell products and services over and above their financial activities (e.g., Mersch 2019). Moreover, large technology companies may increase market concentration by exploiting their network externalities, in turn creating new risks such as the potential for money laundering and abuse of consumer protection and privacy. Facebook in this regard is (again) a recent example. Other risks may especially arise from big tech's funding models, which often use a mix of internal and external investors to finance loans or repackage and sell them to third-party investors (Carstens 2018). In essence, this would be very similar to the originate-to-distribute models that banks conducted in the run up to the financial crisis 2007/2009. As the experience with subprime mortgages has shown, the originate-to-distribute business creates substantial information asymmetries, misaligns incentives and induces financial instability.

In this regard, the German G20 Presidency, as part of its focus on digitalization, has suggested that the Financial Stability Board (FSB) build on the monitoring to date and identify supervisory and regulatory issues of FinTech that merit authorities' attention from a financial stability perspective. Drawing on the findings of the literature, discussions with academics and industry participants, and a stocktake of regulatory approaches to fintech, a number of issues can be highlighted to merit authorities' attention. Particularly, three areas are seen as priority issues, which will require international collaboration (FSB 2017). First, authorities should determine if current

oversight frameworks for important third-party service providers to financial institutions are appropriate, e.g. in cloud computing and data services, in particular, if financial institutions rely on the same third-party service providers. To manage the inherent operational risk efficiently and effectively, greater coordination globally across financial authorities may be necessary, including with non-traditional partners such as authorities responsible for information technology safety and security. Second, ex-ante contingency plans for cyber-attacks, information sharing, monitoring, a focus on incorporating cyber-security in the early design of systems, and financial and technology literacy could help to lower the probability of cyber events that have adverse effects on financial stability. And third, regulators should monitor the FinTech market closely to be able to detect concentration risks and sudden adverse funding flows on FinTech lending platforms with large and unstable implications for broader financial markets.

Other issues that merit authorities' attention include (i) cross-jurisdictional compatibility of national legal frameworks in the context of cross-border lending, trading and payment transactions, including via smart contracts, (ii) governance and disclosure frameworks for big data analytics to ensure reduced complexity and opacity of models designed to identify and assess risk, (iii) proper communication channels with the private sector to facilitate exchange of knowledge and expertise, and (iv) building staff capacity in new areas of required expertise. At the same time, regulators and legislative bodies need to ensure that money laundering is prevented, and consumer protection and privacy is guaranteed. In addressing these priority areas, authorities can promote financial stability, fostering responsible innovation and preventing any derailment of authorities' efforts to achieve a more inclusive financial system.

7 Conclusion

The growing body of research on peer-to-peer lending suggests that online credit markets face at least similar problems like traditional lending markets, most of which are related to information asymmetries. While offline credit market relies heavily on verifiable information, notably credit scores and debt-to-income ratios, online credit markets often lack this type of verifiable credit risk assessment. Consistent with Hoff and Stiglitz (1990), Udry (1994), and La Ferrara (2003), evidence shows that in the absence of verifiable information about the borrower, some social networks can help mitigating the

information inequality, in part even compensating the lack of hard borrower risk information (Liberti and Petersen 2018; Lin 2009).

By leveraging the social online network, peer-to-peer lending websites can reduce the cost of individual evaluation of each borrower, thereby circumventing the problem associated with the distributive aspect of access to superior information (Hirshleifer 1971). Evidence from peer-to-peer lending suggests that these online markets may alleviate the problem of credit rationing, and hence may allow for wider access to capital (Lin 2009).

Despite the advancement in information and communication technologies, the fundamental issue of monitoring still exists: it will be inefficient for each investor to monitor the borrower, and whoever is delegated to monitor should assure the investors that he or she is indeed monitoring (again, the ‘monitoring of the monitor’ problem). In the future, new business models could arise to provide a crowd-production solution to this issue. Currently, however, most peer-to-peer lending websites are responsible for the ex-post monitoring, with the assistance from credit bureaus and collection agencies.

While there has been substantial interest in online credit markets, the key question that still remains to be answered is whether they could be a viable alternative to providing credit to segments of the population that do not have access to credit from the formal banking system. One of the main challenges still remains asymmetric information, as it is difficult to assess credit worthiness for segments of the population that do not have formal financial data. Despite the initial promise, most of the online credit markets primarily rely on data provided from credit rating agencies. The few exceptions that have significantly leveraged other sources of formal data are Ant Financial (Alibaba) and Webank, to name a few. Although these online credit markets hold the promise of alleviating credit constraints for the section of population that do not have access to the formal banking system, the question that still lingers is whether they would deliver on this promise.

References

- Adams, W., L. Einav, and J. Levin. 2009. Liquidity constraints and imperfect information in subprime lending. *American Economic Review* 99 (1): 49–84.
- Akerlof, G.A. 1970. The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84 (3): 488–500.
- Ang, J., J.W. Lyn, and F. Tyler. 1995. Evidence on the lack of separation between business and personal risk among small businesses. *Journal of Small Business Finance* 4 (2): 197–210.

- Arrow, K.J. 1973. Higher education as a filter. *Journal of Public Economics* 2 (3): 193–216.
- Ashta, A., and D. Assadi 2009. An analysis of European online micro-lending websites. *Cahiers du CEREN* 29: 147–160.
- Avery, C. 1998. Strategic jump bidding in English auctions. *Review of Economic Studies* 65 (2): 185–210.
- Bachmann, A., A. Becker, D. Buerckner, M. Hilker, F. Kock, M. Lehmann, P. Tiburtius, and B. Funk. 2011. Online peer-to-peer lending—A literature. *Journal of Internet Banking and Commerce* 16 (2).
- Balyuk, T. 2018. *Financial innovation and borrowers: Evidence from peer-to-peer lending*. Working Paper.
- Becker, G.S. 1957. *The economics of discrimination*. Chicago: University of Chicago Press.
- Berger, A.N., N.H. Miller, M.A. Petersen, R.G. Rajan, and J.C. Stein. 2005. Does function follow organizational form? Evidence from the lednign practices of large and small banks. *Journal of Financial Economics* 76: 237–269.
- Blackwell, D.W., and D.B. Winters. 1997. Banking relationships and the effect of monitoring on loan pricing. *Journal of Financial Research* 20: 275–289.
- Boot, A.W.A. 1999. European lessons on consolidation in banking. *Journal of Banking and Finance* 23: 609–614.
- Boot, A.W.A., and A.V. Thakor. 2000. Can relationship banking survive competition? *Journal of Finance* 55 (2): 679–713.
- Buchak G., G. Matvos, T. Piskorski, and A. Seru. 2018. Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130 (3): 453–483.
- Calomiris, C., C.M. Kahn, and S.D. Longhofer. 1994. Housing finance intervention and private incentives: Helping minorities and the poor. *Journal of Money, Credit, and Banking* 26 (3): 634–674.
- Carstens, A. 2018. *Big tech in finance and new challenges for public policy*. Keynote Address at the FT Banking Summit, London, December.
- De Roure, C., L. Pelizzon, and A. Thakor. 2018. *P2P lenders versus banks: Cream skimming or bottom fishing?* Working Paper.
- Diamond, D. 1984. Financial intermediation and delegated monitoring. *Review of Economic Studies* 51 (3): 393–414.
- Duarte, J., S. Siegel, and L. Young. 2012. Trust and credit: The role of appearance in peer-to-peer lending. *Review of Financial Studies* 25: 2455–2484.
- Financial Stability Board. 2017. *Financial stability implications from fintech*, available at <https://www.fsb.org/2017/06/financial-stability-implications-from-fintech>.
- Freedman, S., and G.Z. Jin. 2011. *Learning by doing with asymmetric information: Evidence from prosper.com*. National Bureau of Economic Research Working Paper 16855.
- Freedman, S., and G.Z. Jin. 2017. The information value of online social networks: Lessons from peer-to-peer lending. *International Journal of Industrial Organization* 51: 185–222.

- Fuster, A., M. Plosser, P. Schnabl, and J. Vickery. 2018. *The role of technology in mortgage lending*. National Bureau of Economic Research Working Paper 24500.
- Gorton, Gary, and Andrew Winton. 2003. Financial intermediation. Handbook of the economics of finance. In *Handbook of the Economics of Finance*, Ed. 1, Volume 1, Chapter 8, ed. G.M. Constantinides, M. Harris and R. M. Stulz, 431–552. Elsevier.
- Grossman, S.J., and J.E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70 (3): 393–408.
- Herzenstein, M., S. Sonenshein, and U.M. Dholakia. 2011. Tell me a good story and i may lend you money: The role of narratives in peer-to-peer lending decisions. *Journal of Marketing Research* 48: 138–149.
- Hildebrand, T., M. Puri, and J. Rocholl. 2016. Adverse incentives in crowdfunding. *Management Science* 63: 587–608.
- Hirshleifer, J. 1971. The private and social value of information and the reward to inventive activity. *American Economic Review* 61 (4): 561–574.
- Hoff, K., and J. Stiglitz. 1990. Imperfect information and rural credit markets—Puzzles and policy perspectives. *World Bank Economic Review* 4 (3): 235–250.
- Iyer, R., A.I. Khwaja, E.F. Luttmer, and K. Shue. 2016. Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62: 1554–1577.
- La Ferrara, E. 2003. Kin groups and reciprocity: A model of credit transactions in Ghana. *American Economic Review* 93 (5): 1730–1751.
- Liao, L., Z. Wang, H. Xiang, and X. Zhang. 2017. *P2P lending in China: An overview*. Working Paper.
- Liberti, J.M., and M.A. Petersen. 2018. *Information: Hard and soft, national bureau of economic research*. Working Paper 25075.
- Lin, M. 2009. Peer-to-peer lending: An empirical study. In *Proceedings of the Americas Conference on Information Systems (AMCIS)*.
- Lin, M., N.R. Prabhala, and S. Viswanathan. 2013. Judging borrowers by the company they keep: Friendship networks and information asymmetry in online peer-to-peer lending. *Management Science* 59: 17–35.
- Lundberg, S.J., and R. Startz. 1998. *Inequality and race: Models and policy*. Discussion Papers in Economics at the University of Washington 0067.
- Manove, M., J. Padilla, and M. Pagano. 2001. Collateral vs. project screening: A model of lazy banks. *RAND Journal of Economics* 32 (4): 726–744.
- Mersch, Y. 2019. *Lending and payment systems in upheaval: The fintech challenge*. Speech given at the 3rd annual Conference on Fintech and Digital Innovation, 26 February 2019, Brussels.
- Michels, J. 2012. Do unverifiable disclosures matter? Evidence from peer-to-peer lending. *The Accounting Review* 87: 1385–413.
- Phelps, E. 1972. The statistical theory of racism and sexism. *American Economic Review* 62 (4): 659–661.
- Pope, D., and J. Sydnor. 2011. What's in a picture? Evidence of discrimination from prosper.com. *Journal of Human Resources* 46: 53–92.

- P2P Banking. *P2P lending marketplace news and reviews*. Available at: <https://www.p2p-banking.com/>.
- Rau, P.R. 2019. *Law, trust, and the development of crowdfunding*. Working Paper.
- Ravina, E. 2018. *Love & loans: The effect of beauty and personal characteristics in credit markets*. Working Paper.
- Spence, M. 1973. Job market signaling. *Quarterly Journal of Economics* 87: 355–374.
- Stiglitz, J.E., and A. Weiss. 1981. Credit rationing in markets with imperfect information. *American Economic Review* 71: 393–410.
- Tang, H. 2018. Peer-to-peer lenders versus banks: Substitutes or complements? *Review of Financial Studies*, forthcoming.
- Udry, C. 1994. Risk and insurance in a rural credit market: An empirical investigation in Northern Nigeria. *Review of Economic Studies* 61 (3): 495–526.
- Williamson, S.D. 1987. Costly monitoring, Loan contracts, and equilibrium credit rationing. *Quarterly Journal of Economics* 102 (1): 135–145.
- Wolfe, B., and W. Yoo. 2017. *Crowding out banks: Credit substitution by peer-to-peer lending*. Working Paper.



Equity Crowdfunding and Capital Formation for Early-Stage Firms

Armin Schwienbacher

1 Introduction

Access to financial resources is fundamental to early-stage firms, as they typically need to fund their innovation activities with external resources (Cosh et al. 2009). A major obstacle is that early-stage firms are particularly difficult to finance, since they face severe agency costs and information asymmetry (Berger and Udell 1998; Carpenter and Petersen 2002). Banks generally do not finance early-stage firms, whose financing is more easily solved with equity than debt. Traditionally, only venture capital funds and business angels have offered (specialized) finance at this stage of development, although most often insufficiently, leading to a funding gap (Cressy 2002). More recently, new forms of early-stage financing have emerged through the digitalization of finance, one of which is crowdfunding.

A. Schwienbacher (✉)
SKEMA Business School, Lille, France
e-mail: armin.schwienbacher@skema.edu

Equity crowdfunding, as we know it today, started in 2008–2009 with the arrival of online platforms that act as intermediary between entrepreneurs and crowd investors.¹ Earlier attempts of raising equity finance online are rare. One such example from the United Kingdom is Trampoline Systems in 2009, which managed to raise the capital (£1 million) without a platform. However, given the regulatory uncertainty at that time, participation in the funding was eventually restricted to accredited investors (Leboeuf and Schwienbacher 2018). Today, a large number of equity crowdfunding platforms coexist and enable the participation of retail investors.

The development of equity crowdfunding was facilitated by regulatory initiatives around the world (Rau 2018) that either made it legal or—when it was already legal, such as the European Union—reduced legal uncertainty on how to run a campaign (Cumming and Johan 2013; Hornuf and Schwienbacher 2016, 2017). Equity crowdfunding received also significant media attention through the JOBS Act in the United States that was promoted by President Obama. Since equity crowdfunding involves issuing financial securities, it falls under securities regulation, an area that is heavily regulated for general solicitations. Most of the campaigns therefore operate under specific exemptions to avoid complying with the costly requirements of a formal prospectus. Exemptions often exist, but involve restrictions in how much can be raised or how many investors can be contacted. In recent years, many countries have clarified and expanded the set of exemptions to facilitate equity crowdfunding and enabling larger issuances. In the European Union, for instance, it is now possible in different member states—under certain conditions—to raise up to EUR 5 million without issuing a formal prospectus.

As with all segments of fintech (crowdfunding is one segment of fintech), an important question often raised is whether equity crowdfunding will disrupt existing activities of banks and other financial intermediaries. At the present time, this appears unlikely to be happening. However, this question is not the most fundamental one, in particular from the perspective of entrepreneurial finance. There, the question is whether equity crowdfunding solves a funding gap or market inefficiency, regardless of its impact on existing sources of finance. Different ideas have been suggested on how this may occur. For instance, Agrawal et al. (2016) offer an interesting discussion on

¹ Other types of crowdfunding exist, such as reward-based, loan-based, and donation-based crowdfunding. They generally involve different types of projects and attract different types of backers, notably through the type of compensation they offer to the backers. Many of them are also project-based. Our focus is on early-stage capital formation and thus is restricted to equity crowdfunding that involves necessarily to legal entity (a firm) that offers securities and thus offers financial returns. See, for instance, Schwienbacher (2013) for more details on these other types of crowdfunding.

how equity crowdfunding is likely to affect how venture capital will operate in the future. Syndication may provide the right benefits for equity crowdfunding by offering more resources while at the same time letting informed investors lead and do the monitoring of startups. This would allow equity crowdfunding to overcome many of its drawbacks, including the fact that they are widely dispersed and often lack knowledge for proper deal selection and monitoring post-investment. According to statistics provided by Agrawal et al. (2016), syndicated deals have started to become the main form of deal funding on the Angel List platform starting in 2014. Although this platform is very specific, it may represent a business model that other platforms can adopt in the future. Catalini et al. (2016) discuss how equity crowdfunding can facilitate access to finance for discriminated people as a way to democratize access to financial resources. If crowdfunding can achieve that goal, it would have a clear societal reason to exist. Similarly, Sorenson et al. (2016) conclude that crowdfunding helps expand the financing of innovative projects beyond the traditional areas financed by venture capital funds. Crowdfunding is therefore less concentrated than venture capital in the United States, for which 50% of the capital goes to four counties only (Sorenson et al. 2016). More importantly, they argue crowdfunding attracts venture capital in new areas as some of initially crowdfunded projects in remote areas are able to attract the attention of venture capital funds. Most of these results are based on reward-based crowdfunding however, and many of these benefits still need to be shown for equity crowdfunding.

Equity crowdfunding has witnessed significant growth rates over the past few years (Rau 2018), similar to other types of crowdfunding (though the magnitude in terms of market size is much smaller). All types combined, crowdfunding generated a worldwide volume of USD 290 billion in 2016, while it barely generate a volume of USD 0.5 billion in 2011. In developed countries, 16% of this total volume is equity crowdfunding (Rau 2018).

This chapter's objective is to provide an overview of the important streams of literature that have emerged in equity crowdfunding research. The bulk of it is empirical in nature due to the increasing availability of data that have either been hand-collected by researchers or obtained from platforms directly. The remainder of this chapter is structured as follows. Section 2 presents details on the functioning of platforms and its peculiarities compared to traditional financing sources. Section 3 discusses the type of firms launching equity crowdfunding campaigns and success drivers. Section 4 elaborates on the range of securities offered and contractual arrangements seen in the investment documentation. Section 5 concludes.

2 Operating Models of Equity Crowdfunding Platforms

As mentioned earlier, equity crowdfunding campaigns launched by early-stage firms take place nowadays on online platforms that have standardized the process and the offerings (Hornuf and Schwienbacher 2016). These platforms intermediate between entrepreneurs and crowd investors. In contrast to reward-based crowdfunding platforms where almost anyone can post a project subject to light selection criteria (mainly ethical criteria and sometimes restricted to specific categories when the platform is specialized), equity crowdfunding platforms perform thorough selection, so that in many cases less than 5% of the entrepreneurs get their project posted. Although there is no empirical evidence yet, it seems reasonable to claim that platforms try to propose projects that are similar in terms of industry focus to what venture capital funds also invest. However, they tend to be smaller in scale.

A crucial element of equity crowdfunding is to enable securities issuance at lower costs (by relying on exemptions in the securities and prospectus regulations) at the expense of lower investor protection compared to securities issuances done by large firms. Also, typical intermediaries and certifiers such as investment banks and underwriters do not participate in order to reduce costs. However, this neglects indirect costs, since crowd investors need to collect and process the information and do much of the due diligence themselves, and cannot rely on intermediaries who perform these tasks for investors.

Some platforms such as WiSEED in France even incorporate an e-vote phase where members of the platform can vote and provide feedback to the platform about their investment interest in the project (Cumming et al. 2019). Once an entrepreneur has successfully passed the due diligence stage, a negotiation takes place between him/her and the platform about the investment terms (the 'contract') that will be proposed to crowd investors during the campaign. Platforms are then remunerated mostly by a success fee (typically a percentage of the total amount raised) in case the campaign is successful. Cumming et al. (2019) however question whether getting the crowd involved in the due diligence process is an effective way of choosing the right startups. WiSEED, crowd investors could rate projects along 11 dimensions and report (but no firm commitment) how much they would invest if the project gets to do a campaign. They find that grades are unrelated to campaign outcome, and that voters only invest €18 for every euro promised. These results suggest that the information collected by the platform at the due diligence phase is not particularly useful. However, they also show that

the aggregate amount of capital committed by voters is a strong predictor of campaign outcome. This is a surprising result that hints to the idea that collectively votes cast are informative.

Platforms operate as two-sided markets that need to bring together entrepreneurs in need of financing and crowd investors who are looking for investment opportunities. At the same time, they operate in a financial environment that is populated by other types of financial intermediaries, including venture capital firms and business angels. Thus, platforms need to offer a competitive advantage over these other financial intermediaries to sustain in the long run. Therefore, beyond offering financial returns to crowd investors that should compensate them for the risk they take, equity crowdfunding also needs to either fill a gap or be superior to other financing forms in order to be viable. If not, entrepreneurs may not opt for that form of financing.

One advantage often heard for crowdfunding is that it relies on the 'wisdom of the crowd'. This argument is based on studies dating back to Galton (1907) on the phenomenon of 'vox populi'. Sociologists have extensively studied crowd interactions and identified conditions under which a crowd may be wiser than a single individual (Larrick et al. 2011). Whether this phenomenon is at play in equity crowdfunding remains to be shown. Mollick and Nanda (2015) show that decisions made by the crowd are not that different from those of experts in the field of arts in projects posted on reward-based crowdfunding platforms. This suggests both have similar 'tastes' or able to predict well which art projects are likely to attract public interest. In equity crowdfunding, the phenomenon needs to be tested separately, since crowd investors are not consumers but make investment decisions in firms with innovative projects. These may require specialized skills to assess, which crowd investors may not have. These skills will differ between experts (e.g., a venture capital fund manager) and crowd investors. However, Signori and Vismara (2018) have shown that entrepreneurial firms that raise funds successfully on an equity crowdfunding platform tend to be also successful afterwards in raising further funds. This suggests crowd investors pick good projects on average, although raising further funds does not necessarily lead to significant financial returns. Also, an interesting follow-up question from the wisdom of the crowd phenomenon is whether these investments decisions are better than the opinion of an expert. More research is needed to understand the true value-added of the crowd beyond providing financial resources to these startups.

A peculiarity of crowdfunding compared to other sources of funding is its impact on operating risk. In bank or venture capital financing, investors and

entrepreneurs take substantial risk, since they first invest in productive assets, whose costs are then sunk. If it turns out there is no sufficient market demand for the innovative product, it will lead to losses for the entrepreneur. Also, the bank is likely to be unable to recover the loan. In crowdfunding, crowd investors may also be consumers so that having a successful campaign may signal high demand. This signal is obtained before any investment is sunk. Thus, operating risk is lower (Schwienbacher, 2018). This is particularly the case for reward-based crowdfunding, but can also apply under certain conditions for equity crowdfunding. Indeed, Colombo and Shafi (2016) argue that the quality of the equity crowdfunding campaign gives an indication of the quality of management, since the latter is required to successfully raise large funds from the crowd and the call for financing takes place in the public domain. This leads to information disclosure on managerial skills. While the study by Colombo and Shafi (2016) focuses on reward-based crowdfunding, the aspect of managerial quality seems also applicable for equity crowdfunding.

Platforms typically adopt one of two funding models. One is the all-or-nothing model, which is the most widely used. There, the entrepreneur sets a funding goal, which then becomes binding. If the entrepreneur is not able to attract sufficient funds to reach that goal, funding is cancelled and the entrepreneur does not get anything. The other is the keep-it-all model, which is—for instance—used in some of the platforms in France. There, the entrepreneur indicates a desired funding goal, which, however, is not binding and therefore merely indicative of what they need. The entrepreneur gets whatever is pledged, even when the funding goal is not achieved. Hakenes and Schlegel (2014) argue theoretically that the first model helps attract more funds, since the entrepreneur commits to return the money pledged if not enough crowd investors are convinced of the firm's project. This induces individuals to pledge early on to signal interest, since they do not bear the risk of financing an under-funded firm. This eventually leads firms to raise more funds under the all-or-nothing model. In other words, opting for the all-or-nothing funding model is a costly signal, consistent with Spence (1973). It is costly for the entrepreneur, since he/she then takes the risk of getting nothing if the funding goal is not achieved. Cumming et al. (2020) provide empirical support for this prediction for reward-based crowdfunding. So far, no analysis exists for equity crowdfunding, since such analysis requires investigating a platform that offers the choice between the two funding models. However, the results from Cumming et al. (2020) can be transposed to equity crowdfunding, which involves similar risk bearing scenarios as in reward-based crowdfunding when startups are underfunded.

Consistent with this funding model argument, one can also observe distinct patterns of funding dynamics during campaigns. Vismara (2018) finds evidence of informational cascades in equity crowdfunding so that the participation of more crowd investors in the first days of the campaign attracts more crowd investors in the remaining days of the campaign. Thus, the first days of the campaign are crucial for entrepreneurs, even if they do not obtain the full amount at the very beginning. However, his/her capacity to attract early on crowd investor is important to generate a positive dynamics. The underlying channel of herding may be attributed to rational or irrational behaviour, depending on how the information of the 'hype' is used by crowd investors. Hornuf and Schwiendbacher (2018) show that the within-campaign funding dynamics also depend on how the platform sets up the allocation of shares to investors. Under a first-come-first-served mechanism, a large proportion of pledges come early on, and the number of additional pledges decays quickly. The reason is that crowd investors have no incentives to delay their decision since the price is fixed and there is a risk of not being able to acquire any shares when waiting too long, since there is a limited number of shares offered. Under the information cascade argument of Vismara (2018), this provides valuable feedback to the entrepreneur during the first days of the campaign about the ultimate success chances of the campaign. Under a price bidding mechanism, many investors will wait until the end to participate since they do not want to signal their interest in the project (since otherwise they may contribute to the informational cascade). If they do, they may drive the ultimate price of shares to increase above the minimum price. By bidding late, they avoid counter-bids that would lead to price increases. These findings support the claim that (at least some) crowd investors act strategically, at the expense of not providing indication of interest to the entrepreneur early on during the campaign.

3 Type of Firms Seeking Equity Crowdfunding

Several studies offer insights into drivers of campaign success. This strand of literature is helpful in understanding which firms should consider equity crowdfunding as a source of capital. One of the first study, if not the first one, on equity crowdfunding is the one by Ahlers et al. (2015). It studies factors affecting campaign success and shows that traditional signals in entrepreneurial finance (e.g., awards that can act as external certification) also work in equity crowdfunding where investors are not able to perform the

same extent of due diligence as professional investors (venture capital funds and business angels). Mohammadi and Shafi (2017) investigate whether female investors attract more female investors to participate in an equity crowdfunding campaign, but find no evidence of such type of herding. Hervé et al. (2019) find that location of crowd investors matters. Those located in more 'sociable' areas tend to invest more. These findings are consistent with other studied in finance (Grinblatt and Keloharju 2001; Hong et al. 2004; Ivković and Weisbenner 2007; Kaustia and Knüpfer 2012), showing that social interactions affect participation in stock markets and investment decisions. For instance, Ivković and Weisbenner (2007) find that individuals spending more time interacting with their neighbours are more likely to make similar stock purchases as them. Sociability is considered an important source of communication and information collection through informal ties. Vismara (2016) offers evidence on the impact of equity retention and social network as a means to increase campaign success. Vismara interprets the positive impact of equity retention on campaign success as evidence of signalling effects by entrepreneurs who retain more. More research is expected in this area as new data becomes available.

A common view is that firms launching an equity crowdfunding campaign have limited financing needs as compared to those seeking VC finance, since the typical amounts that are raised are too small for venture capital funds. Recent studies show there are complementarities between crowdfunding and other forms of financing, although such complementarity has been shown especially for reward-based crowdfunding so far. In particular, recent studies have further shown that a successful crowdfunding campaign may be useful to attract venture capital as follow-up funding (Colombo and Shafi 2016; Drover et al. 2015), supporting the idea that completing a successful campaign provides a strong signal that the entrepreneur has potential and is able to interact with the market. This may ultimately reduce uncertainty and information asymmetry through certification (Colombo and Shafi 2016). While their study uses data from reward-based crowdfunding, it seems reasonable to assume that the same conclusions may apply to equity crowdfunding, although in a different way. Indeed, similar to the process where securing the participation of many backers in a reward-based crowdfunding campaign may signal good market potential, the same in an equity crowdfunding campaign could signal acceptance of the entrepreneurial idea and the capacity of the entrepreneur to manage the firm well. This may affect the perception that professional investors will have of the crowdfunded firm. One recent study supporting this view is the one by Hornuf et al. (2018). Using a sample of German and UK equity crowdfunding sample,

they show that startups that were most successful in their equity crowdfunding campaign were also more likely to raise follow-up funding from venture capital or business angel. Similar to Colombo and Shafi (2016), this is an important result on the possible complementarity of equity crowdfunding and other professional equity investors.

Compared to reward-based crowdfunding, Belleflamme et al. (2014) show that startups with larger financing needs are more likely to opt for equity crowdfunding, since reward-based crowdfunding constitutes a form of price discrimination between backers participating in the crowdfunding campaign and those in the aftermarket. Friends, family and fans may be willing to participate in the campaign early on and pay a different price for the reward than the ultimate price of the product in the aftermarket, leading to two different prices. Often the price in the aftermarket is lower than during the campaign. In other times, entrepreneurs will need to offer a discount during the campaign to raise enough funds to start the project, leading to inefficient pricing. Startups with larger financing needs are required to price discriminate more to attract enough individuals during the campaign instead of letting them participate in the aftermarket, which eventually becomes inefficient in the event of excessive price distortion. In this case, reward-based crowdfunding becomes inefficient. Therefore, firms with larger financing needs will opt for equity crowdfunding. There, entrepreneurs offer a share in the firm's value to backers, but cannot price discriminate anymore.

A widely discussed subject, among legal scholars in particular, is whether the equity crowdfunding market is similar to a market for lemons in that only firms that have been rejected by venture capital firms and banks end up on platforms (see, e.g., Ibrahim 2016). For instance, Griffin (2013) raises concerns that entrepreneurs may exploit the lack of crowdfunding investor knowledge in this type of investments to raise funds for inefficient projects. Hazen (2012) therefore argues that proper disclosure is needed to ensure that equity crowdfunding does not become 'fraudfunding'.

Other studies claim the firms seeking equity funding from the crowd do not have to be lemons, since there are valid reasons for firms to prefer equity crowdfunding over other financing sources. One is that a crowdfunding campaign offers great visibility to the firm, something that can be valuable to consumer-oriented firms. Raising venture capital or taking a loan does not provide the same level of visibility. Another reason is that many early-stage firms do simply not aim at the high growth rates that are required for raising venture capital. Indeed, venture capital is very selective so that 99% of the firms that start every year do not even qualify because of lack of growth opportunities (Metrick 2006).

Little empirical evidence is available that enables to answer the ‘lemons’ question. One reason is that most campaigns were run recently and thus it is too early to infer anything about their profitability. Also, firms are unlikely to voluntarily disclose whether they seek equity crowdfunding because they were turned down by traditional investors and banks. One way to infer it—at least partially—is to analyse the firms’ capital structure at the time of their campaign, and examine whether they are already (excessively) leveraged. This view is consistent with the pecking order theory that firms only raise equity finance when their debt capacity is exhausted. However, it does not allow us to infer that equity crowdfunded startups are lemons, but simply that equity crowdfunding is used as last resort; i.e., it is not their preferred mean of financing. This approach is taken by Walthoff-Borm et al. (2018), who show that the average startup on Crowdcube is already highly levered and has little internal resources. Their findings support the idea that the choice of equity crowdfunding as a source of financing for early-stage firms is consistent with the pecking order theory.

4 Contractual Arrangements and Securities Issued

As mentioned before, the terms offered to the crowd are mentioned in a purchase contract that the platform drafts together with the entrepreneur (Hornuf and Schwienbacher 2016). This investment documentation is posted online during the campaign period so that members of the platform can scrutinize it. Next to the price and type of securities, it details the rights of shareholders.

There is a common understanding that one benefit of equity crowdfunding for the entrepreneur is that he or she does not need to hand over as much control rights as under venture capital and business angel finance. The underlying reasons are that (i) the crowd is relatively passive so that they are unlikely to enforce control rights in the first place; and (ii) crowd investors are dispersed, with each holding only a small equity stake. This in turn reduces each crowd investor’s incentives to intervene due to the free-rider problem (Hornuf and Schwienbacher 2016).

However, recent work by Klöhn et al. (2018) and Hornuf et al. (2021) on Germany shows a more nuanced view. Based on a sample of contracts offered on different German equity crowdfunding platforms, they show that contracts are not that different from those signed by venture capitalists (Cumming 2008; Kaplan and Stromberg 2003), as similar cash flow and

control rights are found in both contract types. This of course does not mean they are necessarily enforced in equity crowdfunding. Hornuf et al. (2018) find that crowd investors pay higher prices for securities that offer more control rights, consistent with control right theory (Aghion and Bolton 1992). These conclusions are in line with related studies on venture capital contracts (Cumming 2008; Kaplan and Stromberg 2003), suggesting contract design and pricing can be explained in similar ways as for venture capital investments despite differences in ownership concentration. Thus, while the question of enforceability remains to be examined, their work shows that many of the contractual arrangements and covenants found in venture capital contracts can also be found in equity crowdfunding contracts.

Most of the platforms in Germany use participating notes (Klöhn et al. 2018; Hornuf et al. 2018). Other platforms such as Crowdcube offer common shares to investors. This allows them to participate more actively in the decision process differently than with participating notes. Also, the control rights become less contractual, since investors then hold direct ownership in the startups. Cumming et al. (2018) study the UK-based platform Crowdcube that offers dual-class shares to investigate investors' preferences for voting rights in equity crowdfunding. In particular, entrepreneurs can sell voting shares at a higher price than non-voting shares. They find, among other things, that professional investors always purchase voting shares, but that a too big separation of voting and non-voting shares lowers the chances of a successful crowdfunding campaign. This result is consistent with the view that this outcome leads to higher separation of ownership and cash flow rights among investors.

Another important point noting is that many platforms operate differently than in Germany, as documented by Hornuf et al. (2018). Some pool crowd investors into a financial vehicle that collects the money and then invests the funds in the startup. Crowd investors do not invest directly into the startups, but indirectly through the financial vehicle. The latter then exercises the control rights on behalf of crowd investors. This investment scheme has become particularly common in France (Hervé et al. 2019), where the platform typically acts as trustee in the vehicle in exchange for management fees. In this case, the free-riding problem is largely solved, but somewhat disconnects the entrepreneur with the crowd.

Other platforms propose different solutions, including requiring a higher minimum investment ticket (Hornuf and Schwienbacher 2018). This then restricts the type of participants to wealthier individuals, perhaps even business angels. Some platforms, such as Anaxago in France, have developed a business angel-type of investor network with minimum investments of

EUR 3000 to EUR 5000. In contrast, other platforms allow participation at much lower amounts, such as EUR 100 or even EUR 5. The choice of the minimum ticket directly affects the type of individuals attracted to the platform and the number of investors participating in a deal. With smaller tickets, the number of crowd investors becomes inevitably much larger, which in turn requires a different investment structure than for an investment that involves a few dozens of wealthier investors.

Finally, recent work has investigated other types of securities offered during equity crowdfunding campaigns, including SAFE in the United States (Green and Coyle 2016; Wroldsen 2017). As equity crowdfunding is expanding and rapidly developing in more countries, research in that area should expand significantly in the coming years. While the first studies in that area are being carried out by legal scholars, there is certainly scope for finance scholars to contribute empirically once more data becomes available.

5 Concluding Remarks

Equity crowdfunding has evolved quickly as a possible alternative source of finance for early-stage firms. While many of those firms have successfully raised funds from the crowd (some through multiple rounds), ultimately the viability of equity crowdfunding lies in the successful delivery of realistic returns to crowd investors. It is still too early to provide evidence on this, since most of the campaigns have taken place only recently. However, analyses on the success of crowdfunded firms will provide important additional evidence on the viability (or not) of equity crowdfunding for the long run.

New technological and regulatory developments may help increase the likelihood for equity crowdfunding to sustain growth, notably the upcoming use of blockchain technology for registering privately held shares. In France, for instance, a new regulation was adopted that authorizes the use of blockchain technology for privately held shares (Decree No. 2017–1674 of December 8, 2017; in French: *Ordonnance n°2017–1674 du 8 décembre 2017 relative à l'utilisation d'un dispositif d'enregistrement électronique partagé pour la représentation et la transmission de titres financiers*). This will reduce transaction costs for selling securities of small businesses and could thus foster the development of secondary markets, a key ingredient to private financing that is still missing to make equity crowdfunding viable in the long run. Another example is regulatory developments to create a European market, since equity crowdfunding markets are currently highly segmented by country. Cross-border investments remain limited in Europe today.

In addition, given the large number of extant platforms the market is likely to consolidate in the future. While entry allows for experimentation with different business models, over time, as platforms observe the most efficient funding structures, it is plausible that their business models may converge. One particular trend that may be emerging is based on increasing collaboration with traditional institutional investors. Business angels, venture capital firms and even banks have started to collaborate more closely with platforms and fund their growth. A greater collaboration with traditional investors may further ensure that weaknesses of independent platforms are mitigated.

Finally, future research should focus more on testing economic theories than explaining the functioning of equity crowdfunding itself. Indeed, since the issuance process takes place entirely online, much more can be measured on individual behaviour that could help shed light into how investment decisions for early-stage firms are made.

References

- Agrawal, A., C. Catalini, and A. Goldfarb. 2016. Are syndicates the killer app of equity crowdfunding? *California Management Review* 58 (2): 111–124.
- Aghion, P., and P. Bolton. 1992. An incomplete contracts approach to financial contracting. *Review of Economic Studies* 59 (3): 473–494.
- Ahlers, G.K.C., D.J. Cumming, Gunther, C., and D. Schweizer. 2015. Signaling in equity crowdfunding. *Entrepreneurship Theory and Practice* 39 (4): 955–980.
- Belleflamme, P., T. Lambert, and A. Schwienbacher. 2014. Crowdfunding: Tapping the right crowd. *Journal of Business Venturing* 29 (5): 585–609.
- Berger, A.N., and G.F. Udell. 1998. The Economics of small business finance: The roles of private equity and debt markets in the financial growth cycle. *Journal of Banking and Finance* 22 (6–8): 613–673.
- Carpenter, R.E., and B.C. Petersen. 2002. Is the growth of small firms constrained by internal finance? *Review of Economics and Statistics* 84 (2): 298–309.
- Catalini, C., C. Fazio, and F. Murray. 2016. *Can equity crowdfunding democratize access to capital and investment opportunities?* MIT Innovation Initiative Lab for Innovation Science and Policy Report.
- Colombo, M.G., and K. Shafi. 2016. *Does crowdfunding help firms obtain venture capital and angel finance?* ENTFIN Conference, Lyon, July. Available at www.em-lyon.com/minisiteen/ReCEntFin/ENTFIN-Conference-2016/Papers.
- Cosh, A., D.J. Cumming, and A. Hughes. 2009. Outside entrepreneurial capital. *Economic Journal* 119 (540): 1494–1533.
- Cressy, R. 2002. Introduction: funding gaps: A symposium. *Economic Journal*, F1–F16.

- Cumming, D.J. 2008. Contracts and exits in venture capital finance. *Review of Financial Studies* 21 (5): 1947–1982.
- Cumming, D.J., and S. Johan. 2013. Demand driven securities regulation: Evidence from crowdfunding. *Venture Capital: An International Journal of Entrepreneurial Finance* 15: 361–379.
- Cumming, D.J., G. Leboeuf, and A. Schwienbacher. 2020. *Crowdfunding models: Keep-it-all vs. all-or-nothing*. *Financial Management* 49, 331–360.
- Cumming, D.J., F. Hervé, E. Manthé, and A. Schwienbacher. 2019. *Testing-the-waters policy with hypothetical investment: Evidence from equity crowdfunding*. *Entrepreneurship Theory & Practice*, forthcoming.
- Cumming, D.J., M. Meoli, and S. Vismara. 2018. *Investors' choice between cash and voting rights: Evidence from dual-class equity crowdfunding*. Working paper.
- Drover, W., M.S. Wood, and A. Zacharakis. 2015. Attributes of angel and crowd-funded investments as determinants of VC screening decisions. *Entrepreneurship Theory and Practice* 41 (3): 323–347.
- Galton, F. 1907. Vox populi. *Nature* 75: 450–451.
- Green, J.M., and J.F. Coyle. 2016. Crowdfunding and the not-so-safe Safe. 102 *Virginia Law Review Online* 168, UNC Legal Studies Research Paper No. 2830213.
- Griffin, Z.J. 2013. Crowdfunding: Fleecing the American masses. *Case Western Reserve Journal of Law, Technology & the Internet* 375.
- Grinblatt, M., and M. Keloharju. 2001. How distance, language, and culture influence stockholdings and trades. *Journal of Finance* 56: 1053–1073.
- Hakenes, H., and F. Schlegel. 2014. *Exploiting the financial wisdom of the crowd—Crowdfunding as a tool to aggregate vague information*. Working paper. Available at: <https://ssrn.com/abstract=2475025>.
- Hazen, T.L. 2012. Crowdfunding or fraudfunding? Social networks and the securities laws—Why the specially tailored exemption must be conditioned on meaningful disclosure. *North Carolina Law Review* 90: 1735–1807.
- Hervé, F., E. Manthé, A. Sannajust, and A. Schwienbacher. 2019. *Determinants of individual investment decisions in investment-based crowdfunding*. *Journal of Business Finance and Accounting* 46 (5–6), 762–783.
- Hong, H., J.D. Kubik, and J.C. Stein. 2004. Social interaction and stock market participation. *Journal of Finance* 59: 137–163.
- Hornuf, L., and A. Schwienbacher. 2016. Crowdinvesting—Angel investing for the masses? In *Handbook of research on venture capital: Volume 3 business angels*, ed. Hans Landström and Colin Mason, 381–397. Edward Elgar.
- Hornuf, L., and A. Schwienbacher. 2017. Should securities regulation promote equity crowdfunding? *Small Business Economics*, 49 (3), 579–593.
- Hornuf, L., and A. Schwienbacher. 2018. Market mechanisms and funding dynamics in equity crowdfunding. *Journal of Corporate Finance* 50, 556–574.
- Hornuf, L., and A. Schwienbacher. 2018. Internet-based entrepreneurial finance: Lessons from Germany. *California Management Review* 60 (2): 150–175.

- Hornuf, L., T. Schilling, and A. Schwienbacher (2021). *The relevance of investor rights in crowdinvesting*, *Journal of Corporate Finance*, 101927.
- Ibrahim, D.M. 2016. Equity crowdfunding: A market for lemons? *Minnesota Law Review* 100.
- Ivković, Z., and S. Weisbenner. 2007. Information diffusion effects in individual investors' common stock purchases: Covet thy neighbors' investment choices. *Review of Financial Studies* 20 (4): 1327–1357.
- Kaplan, S.N., and P. Strömberg. 2003. Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. *Review of Economic Studies* 70: 281–315.
- Kaustia, M., and S. Knüpfer. 2012. Peer performance and stock market entry. *Journal of Financial Economics* 104 (2): 321–338.
- Klöhn, L., L. Hornuf, and T. Schilling. 2018. Financial contracting in crowd-investing: Lessons from the German market. *German Law Journal* 19 (3): 509–578.
- Larrick, R.P., A.E. Mannes, and J.B. Soll. 2011. The social psychology of the wisdom of crowds. In *Frontiers in social psychology: Social judgment and decision making*, ed. Joachim I. Krueger, 227–242. New York: Psychology Press.
- Leboeuf, G., and A. Schwienbacher. 2018. Crowdfunding as a new financing tool. In *The economics of crowdfunding—Startups, portals and investor behavior*, ed. D. Cumming and L. Hornuf, 11–28. Palgrave Macmillan.
- Metrick, A. 2006. *Venture capital and the finance of innovation*. Wiley.
- Mohammadi, A., and K. Shafi. 2017. Gender differences in the contribution patterns of equity-crowdfunding investors. *Small Business Economics*, forthcoming.
- Mollick, E.R., and R. Nanda. 2015. Wisdom or madness? Comparing crowds with expert evaluation in funding the arts. *Management Science* 62 (6): 1533–1553.
- Rau, P.R. 2018. *Law, trust, and the development of crowdfunding*. Working paper. Available at SSRN: <https://ssrn.com/abstract=2989056>.
- Signori, A., and S. Vismara. 2018. Does success bring success? The post offering lives of equity crowd-funded firms. *Journal of Corporate Finance*, 50(C), 575–591.
- Schwiebacher, A. 2013. Financing the business. In *The Routledge companion to entrepreneurship*, ed. Ted Baker and Friederike Welter, 193–206. Abingdon, Oxon: Routledge.
- Schwiebacher, A. 2018. Entrepreneurial risk-taking in crowdfunding campaigns. *Small Business Economics*, 51 (4), 843–859.
- Sorenson, O., V. Assenova, G.-C. Li, J. Boada, and L. Fleming. 2016. Expand innovation finance via crowdfunding. *Science* 354 (6319): 1526–1528.
- Spence, M. 1973. Job market signaling. *Quarterly Journal of Economics* 87 (3): 355–374. <https://doi.org/10.2307/1882010>.
- Vismara, S. 2016. Equity retention and social network theory in equity crowdfunding. *Small Business Economics* 46 (4): 579–590.
- Vismara, S. 2018. Information cascades among investors in equity crowdfunding. *Entrepreneurship Theory and Practice* 42 (3): 467–497.

- Walthoff-Borm, X., A. Schwienbacher, and T. Vanacker. 2018. Equity crowdfunding: First resort of last resort? *Journal of Business Venturing* 33, 513–533.
- Wroldsen, J. 2017. Crowdfunding investment contracts. *11 Virginia Law & Business Review*, forthcoming.



Growing Pains: The Changing Regulation of Alternative Lending Platforms

Ding Chen, Anil Savio Kavuri, and Alistair Milne

1 Introduction

This chapter examines the emerging regulation of alternative lending platforms around the world, focusing primarily on the US, China, and the UK where alternative lending has developed furthest.¹

Here “alternative lending platforms” refers to the technology-based platforms that allow investors to participate as investors in loan assets, with a direct claim on payment of interest and repayments of principal. The platform itself has no claim on these payments, but instead earns fees for related

¹ A summary version of some material from this chapter is previously published in (Chen et al. 2020), along with a more detailed discussion of the experience of P2P lending of China.

D. Chen

School of Law, University of Sheffield, Sheffield, UK

e-mail: ding.chen@sheffield.ac.uk

A. S. Kavuri

Crawford School of Public Policy, Australian National University, Canberra, Australia

e-mail: a.s.kavuri@lboro.ac.uk

A. S. Kavuri · A. Milne (✉)

School of Business and Economics, Loughborough University, Loughborough, UK

e-mail: a.k.l.milne@lboro.ac.uk

Table 1 Lending volumes: consumer marketplace lending

New loans to consumers	2013		2017	
	\$bn	(%) GDP	\$bn	(%) GDP
China	3.85	0.04	224.43	1.83
US	2.81	0.02	14.66	0.08
UK	0.29	0.01	1.40	0.05
RoW	0.12	0.0	1.99	0.004

Source various reports of the Cambridge center for alternative finance; author's calculations

services including the assessment of credit risk, the matching of investors with borrowers, and the servicing of loans including the collection and allocation of payments of interest and principal. Such platforms are referred to by a variety of terms including “P2P lenders” and “loan-based crowdfunders.” They can be distinguished from other technology-based non-bank “balance sheet lenders,” who instead hold loans on their own balance sheets funded more conventionally using a mixture of debt and equity. The term “marketplace lending” is also often used, especially in the US, to refer collectively to both alternative lending platforms and non-bank balance sheet lenders.²

Alternative lending platforms, together with technology-based non-banking balance sheet lenders, are important alternatives to bank-intermediated lending. They have competitive advantages relative to traditional banks. They are freed from the constraints imposed by the legacy of existing bank systems for approving and managing loans, helping them automate the processes of loan application and loan approval. Online applicants can obtain a credit decision within hours or even minutes. Their systems also support detailed granular disclosure about both loan credit assessment and loan performance and can help reduce the costs of loan servicing. As documented [reference to other chapters of the book], these competitive advantages have, since the first establishment of such platforms in 2005, allowed alternative lending to grow rapidly over the past twelve years.

The first emergence and furthest development of alternative lending has been in the US, China, and the UK and these remain the jurisdictions with much the largest volume of lending (Tables 1, 2).

Regulation has been slow to keep up. The overall policy concern is getting the balance right, between, on the one hand, sufficient regulation and supervision to prevent poor business practice and ensure both

² (Perkins 2018) Figure I provides useful visual illustration of the different lending models. His classification of “marketplace lending” as embracing both alternative lending platforms and balance sheet lenders is useful, but it is not universally accepted.

Table 2 Lending volumes business marketplace lending

New loans to business	2013		2017	
	\$bn	(%) GDP	\$bn	(%) GDP
China	1.44	0.01	97.43	0.80
US	0.34	0.00	1.45	0.01
UK	0.19	0.01	2.04	0.08
RoW	0.11	0.00	1.15	0.003

Source various reports of the Cambridge center for alternative finance; author's calculations

borrowers and investors are adequately served by the new platforms; and on the other, supporting innovation that promotes competition, encourages financial inclusion, and improves portfolio opportunities for investors.

This chapter examines the emerging framework for regulation of alternative lending platforms in these three pioneer jurisdictions—the US, China, and the UK (with briefer assessment of developments elsewhere). The chapter is organized as follows. Sections 2 through 4 review the legal and regulatory framework for alternative lending as it has and continues to develop in the UK, US, and China. Section 5 provides a briefer discussion of the legal and regulatory frameworks in the rest of Europe, North America, and Asia where alternative lending, while also growing fast, is at a much earlier stage of development.

Section 6 concludes with a discussion of the challenges of regulating alternative lending and what can be learned from experience of these three countries for achieving an appropriate balance between regulatory oversight and support for innovation in the mechanisms of lending. We argue that the best outcome will be achieved through a technology-based approach to regulation, using new technology to level the playing field between alternative lenders and banks. We hope and anticipate that the framework of law and regulation will evolve in this direction in the US, China, the UK, and in other countries; but this new industry will not be able to avoid substantial regulatory growing pains.

2 Regulation in the UK

2.1 Overview

The UK saw some of the earliest developments in alternative lending platforms, with the 2005 establishment of the first “peer-to-peer” lender Zopa

offering personal loans followed by the 2010 launch of Funding Circle offering small business loans. More than one hundred platforms have now been started in the UK and there has been rapid growth in total value of supported loans.³ Only a few platforms though operate on any scale, and at end-2017 only three intermediated more than £0.5 bn loans on their platforms: Funding Circle (small business loans £1.6 bn), Zopa (personal loans £1.2 bn) and Ratesetter (both personal and small business, £0.8 bn).⁴

This section reviews the developments in the regulation of UK alternative lending platforms. Three phases can be distinguished. Initially, when still nascent, alternative lending largely fell under the regulatory radar (aside from existing consumer credit regulation). A second phase was the development of an initial framework for regulation in 2014–2016, after the transfer of regulatory responsibility from the Office of Fair Trading to the Financial Services Authority (superseded by the Financial Conduct Authority in 2013). This was given important political support by the 2010–2015 coalition chancellor of the exchequer (finance minister) George Osborne. His March 2014 budget introduced the Innovative Finance Individual Savings Accounts (IFISAs) for alternative lending platforms.⁵ This extended the existing ISA scheme for tax-exempt retail savings, allowing qualifying platforms to offer tax-exempt IFISAs from April 1, 2016, within the overall limits for annual ISA savings.⁶

The third phase of regulation is now underway, a strengthening of the regulatory framework following a “post-implementation” review launched by the FCA in 2016 resulting in the conclusions and recommendations of (FCA 2019). This tightening is in part motivated by the increased perception of potential risks due to rising loan arrears, especially in some platforms focused on lending for property development.⁷

³ <http://www.p2pmoney.co.uk/companies.htm> lists 133 platforms as of early Nov 2018, including 30 that had stopped doing P2P or were no longer in operation.

⁴ <http://p2pfa.info>, Ratesetter data from <https://invest.ratesetter.com/aboutus/statistics>.

⁵ See (HM Revenues & Customs 2015). The objective of this policy is described as “To increase the choice and flexibility available to ISA investors, encourage the growth of peer-to-peer lending and improve competition in the banking sector by diversifying the available sources of finance.”

⁶ Also the Autumn Statement of Dec, 2014 amended the calculation of interest income on investment in loan-based crowdfunding for individual tax liability, from gross interest to interest net of any portfolio loan losses.

⁷ (Neuwirth 2017; Megaw 2018). Project overruns in property development make relatively high payment arrears normal in this lending.

2.2 Regulatory Responsibilities and Borrower Protection

Alternative lending platforms (commonly referred to as P2P lenders in the UK) were originally overseen by the Office of Fair Trading (OFT) using licenses for debt administration and consumer lending. This changed in April 2014, when a new regulated activity of “operating an electronic system in relation to lending,” under the Financial Services and Markets Act 2000 (Regulated Activities) Order 2001, came into force and, at the same time, responsibility for consumer credit regulation transferred from the OFT to FCA.

Changes to the FCA regulatory handbook (FCA 2016a) introduced two forms of “crowdfunding” as regulated activities, both requiring FCA authorization⁸:

- I. Investment-based crowdfunding in unlisted corporate equity, unlisted debt-securities, and equity stakes in property and real estate. The new regulations limited ordinary retail investors investment crowdfunding to a small share, less than ten percent of their net financial assets individuals, or who are investing on the basis of regulated advice;
- II. Loan-based crowdfunding (i.e., P2P lending) including the funding of consumer loans, business loans and loans secured on the property, where the regulatory regime continued that established under the OFT, generally permitting retail investment subject to requirements on risk disclosure.⁹

Crowdfunding platforms are expected to fulfill similar conditions to any other FCA regulated entity, demonstrating adequate resources and a suitable business model and complying with general FCA rules on for example management systems and controls and FCA requirements for preventing financial crime. The FCA was cautious about granting authorizations. As of September 2016 the FCA had granted full authorizations for loan-based crowdfunding platforms to only 12 firms, and was assessing 86 additional applications of which 39 were already operating under interim permission

⁸ (Wardrop and Ziegler 2016; Milne and Parboteeah 2016; FCA 2018) describe the evolution of these arrangements. Other forms of crowdfunding not FCA regulated but not authorized are donation-based crowdfunding, where there is no expectation of a return; and “pre-payment or reward based crowd-funding,” where the return is the offer of a products or services e.g. concert tickets or access to a computer game, <https://www.fca.org.uk/firms/authorisation/when-required/crowdfunding>.

⁹ Platform based investment in business invoice e.g. www.marketinvoice.com is also classified as a loan-based crowdfunding, but not marketable to unsophisticated retail investors.

inherited from the previous regime.¹⁰ Loan-based crowdfunding platforms are also required to report quarterly to the FCA.¹¹

While the main regulatory concerns with alternative lending platforms have been about investor communication and protection, the UK regulations also protect borrowers. Platforms must follow general FCA rules on appropriate conduct of their business, including checks on affordability when dealing with borrowers. UK loan-based platforms intermediating personal loans must comply with the same consumer finance law and regulation as other unsecured lenders.¹² The new proposals (FCA 2019) will include strengthened requirements for ensuring all the FCA rules for borrower protection, including those for mortgage and home finance borrowers, are fully applied to lending financed through crowdfunding platforms as well as to traditional intermediated lending

2.3 Systems and Controls

During its development of the initial regulatory framework for loan-based crowdfunding, the FCA focused on systems and controls, paying particular attention to segregation of client assets (ensuring these are not comingled with platform assets); and arrangements for orderly transfer of business in the event of platform failure. As we describe more fully in the next sub-section (FCA 2019) is introducing a far more detailed regulatory regime. These will include a set of detailed and strengthened rules covering the wind-down of platforms. Here every platform must have detailed policies in place in the form of a “resolution manual” describing in practical detail how a wind-down will be implemented.

The forthcoming regime also introduces much stronger rules on governance, aiming to bring P2P platforms in line with the FCA’s expectations of governance in investment management, i.e., properly fulfilling the fiduciary duties of an agent operating on behalf of a relatively uninformed client. They must have clearly stated risk-management policies. Also, among their operational staffing responsibilities platforms must have independent risk-management, regulatory compliance, and internal audit functions (for smaller platforms these functions can be combined with other responsibilities).

¹⁰ (Davidson 2016). Difficult to update as the FCA does not maintain a published list of authorized platforms.

¹¹ <https://www.fca.org.uk/firms/gabriel/crowdfunding-data-guide>, using an XML or Excel proforma covering their financial position and loan performance, client money held, complaints and details of loans arranged each quarter (there is currently no regulatory reporting of charging arrangements).

¹² Summarized in (FCA 2016b).

Finally, they must demonstrate they are aware of and dealing appropriately with conflicts of interest, for example, opaque fee arrangements, participation in loans or loan markets by staff or family members or any other situation where decisions may potentially benefit a client or staff member at the expense of another client.

While alternative lending is still much too small to be regarded as a systemic financial risk, the FCA also applies a graduated system of capital requirements at 0.2% of the initial £50 million of total value of loaned funds outstanding and at lower rates that fall from 0.15% on the next £200 m; to 0.05% of any lending over £500 mn (limiting the risks of the financial failure of a platform and hence protecting both borrowers and investors). This provides additional protection against the risk of a disorderly platform failure triggered by financial difficulties.

2.4 Investor Protection: Transparency and Risk Communication

A major concern of the FCA, when developing its initial 2014 regulatory framework was ensuring that retail investors were fully aware that loan participations through alternative lending platforms are not at all similar to bank deposits. They are not covered by the Financial Services Compensation Scheme (the UK bank deposit insurance scheme) and projected returns are not guaranteed so it is possible that investors could lose some of their initial investment. All the well-established platforms make clear on their webpages and during the process of investment that they are risky investments.

When establishing this first framework for regulating alternative lending, the FSA was already expressing concerns about risks of investment in crowdfunding platforms, including¹³: difficulties for investors comparing platforms with each other or to compare crowdfunding with other asset classes due to complex and often unclear product offerings; difficulties assessing the risks and returns of investing via a platform; financial promotions that do not always meet the FCA requirement to be “clear, fair and not misleading”; complex structures introducing operational risks and/or conflicts of interest that are not always sufficiently well managed; inadequate or unclear risk disclosure to investors; and the use of provision funds that may obscure the underlying risk to investors or lead them to believe that platforms are providing an implicit guarantee of the loans they facilitate.

¹³ See (FCA 2016c).

These concerns about the transparency and risk communication have intensified in the FCA's "post-implementation" review of their crowdfunding regulations (FCA 2018). This notes that: "We observe some poor business practices" (pg 2) including; and "some platforms not only facilitate lending but actively structure which loans investors are exposed to in order to achieve a target rate of return. As a result, their offering is also marketed differently, with a focus on headline rates" (Pg 4).

As a result, the FCA has now introduced in 2019 a substantially more detailed regime of crowdfunding regulation (FCA 2019). The new approach introduces an extensive number of new, detailed rules. These are though intended to be flexible: "... *there are some areas where platforms will have varying implementation options. These include, for example, simplifying elements of their business model, or amending their processes and systems to comply with the more detailed rules*" (FCA 2018, pg 33).

One group of new rules is on the assessment of borrower credit risk and subsequent risk pricing. Platforms are expected to operate to the same established standards as expected for balance sheet lenders who risk their own capital (though the sophistication of their approaches may still vary), with: developed and tested models for both risk classification and pricing; consideration of affordability in line with FCA consumer credit and mortgage conduct of business rules; and systems to ensure that loans allocated to investors are accurately classified according to the criteria advertised to investors. More sophisticated approaches are expected when a platform offers a target or projected rate of portfolio return (as all the larger UK platforms do).

The second group of proposed rules restricts the marketing of investment in alternative lending platforms, going much further than the current requirement of making clear that these are not insured deposits. Direct marketing will now only be possible for three groups: (i) those who confirm *before* the promotion that their decisions will be on the basis of advice or investment management services from regulated advisors or investment managers; (ii) those certified or self-certified as "sophisticated investors" or as "high net-worth investors"; or (iii) those who certify *before the promotion* that they will not invest more than 10% of their net investible portfolio in equity or loan-based crowdfunding; and in addition for cases (ii) and (iii) the platform must be able to demonstrate compliance with the rather complex FCA rules on appropriateness (COBS 10).

The third set of new rules cover disclosure requirements on platform operations. Again, there are extensive new requirements. The platforms must make

clear to investors the processes they undertake in bringing loans to the platform and matching with investors, including: (i) “the nature and extent of due diligence it undertakes for borrowers”; (ii) the criteria and processes for loan assessment and risk classification; (iii) the process for determining the “price” (in the case of a loan the interest rate) and the construction of investor portfolios; (iv) procedures for dealing with loan arrears and default; (v) the terms on which investors can realize their positions early by sale on a secondary market and how prices in that market are determined including the possibility of risk to their investment from resale; (vi) explanation of wind-up arrangements including naming of any third party “back-up servicer” who would take over loan and investment servicing responsibilities in the event of platform failure and treatment of any uninvested client money.

The fourth set of proposed rules covers investment information, including on the operation of any “contingency” or “provision” funds, i.e., arrangements for pooling investor money in order to make good on future loan losses.

After loans have been matched to investors then there must be ongoing disclosure giving investors access to: (i) up to date information on all loans in which an investor participates, on original loan terms including original assessment of credit risk and affordability and whether there is any security on the loan, on previous and remaining capital and interest payments and on fees paid to platform in relation to that borrower; (ii) update on any change in likely actual returns; (iii) whether there has been default and following a default an estimated value of the P2P agreement (with a standardized definition of default as 90 days past due, except in the case of property loans where it will be 180 days past due).

In addition, where platforms sets a price for loans (as an alternative to loan prices being set in investor auctions), then it must publish “outcomes statements” within four months of the end of each financial year, including expected and actual default rates for all P2P agreements by risk category, a summary of assumptions used to determined expected default rates; and a comparison with any originally advertised prospective “target” rate.

Finally, a further part of the proposed rules, are additional disclosures on “contingency” or “provision” funds: first making absolutely clear that any such fund is *not* a guarantee of return or equivalent to the insurance of bank deposits offered under the Financial Services Compensation Scheme (unlike the FSCS there is no further obligation on the platform or other industry participants to make good losses once the fund is exhausted); and second a requirement for a clear statement to investors on “*How the contingency fund is funded how the contingency fund is governed who the money in the fund belongs to the considerations the fund/platform takes into account when deciding whether*

or how to exercise its discretion to pay out from the fund, including examples. This should include consideration of whether or not the fund has sufficient money to pay and that it has absolute discretion in any event not to pay or to decide the amount of the payment the process for considering pay outs from the fund a description of how that money will be treated in the event of the platform's insolvency” (FCA 2018, pg 51).

3 Regulation in the US

3.1 Overview

This section describes developments in the regulation of US alternative lending platforms, since the launch of Prosper in 2006. Prosper was the first US platform. Two issues have been prominent. First is the way that the business models of the lending platforms have had to adapt to the maze of federal and state regulation developed in the historical context of face-to-face delivery and local competition in banking, investment services, and consumer lending.¹⁴ Compliance with state-based regulation is an ongoing concern for all platforms facilitating online loan investment across state borders. The second issue is the application of SEC investor protection rules under Federal securities law to alternative lending platforms, imposing substantial compliance requirements particularly burdensome on small platforms.

In the US, unlike in the UK, there is no specific regulatory framework for “loan-based crowdfunding” (or P2P) platforms, i.e., those platforms which facilitate direct investment in loans, without holding whole loans or a share of loans themselves on their own balance sheets. Instead, these platforms operate within the extensive regulatory landscape that applies also to a variety of non-bank “balance-sheet” lenders.

From 2006 until 2008, US alternative lending developed with relatively little regulatory oversight other than that of existing banking and consumer lending regulation. Subsequently, two major developments stand out. In 2008 Securities and Exchange Commission (SEC) ruled that the “notes” (claims on payments of loan interest and principal) issued by platforms to investors were not exempt from securities laws and had to be registered with the SEC in the same manner as any other investor securities. In 2015 the SEC

¹⁴ The complexity of this maze is documented in the annual summaries of the US regulatory framework for marketplace lending published by the law firm Chapman and Cutler LLP, of which the latest is (Manbeck and Franson 2018).

issued detailed rules on how these securities laws were to be applied to alternative lending platforms. This regulatory approach, imposing strong protections on the sale of notes to retail investors, has been one of the main reasons why alternative lending in the US is now overwhelmingly an institutional not retail investor asset class.

The second development has been the evolving application of state consumer borrowing protection to alternative lending platforms, in particular state “usury” laws that set maximum interest rates on consumer loans. The 2016 *Madden v Midland Funding LLC*¹⁵ case decision of the Second Circuit called into question a core established practice of US alternative lending platforms: that of originating their loans through a national (Office of the Comptroller of the Currency regulated) bank or Federal Deposit Insurance Corporation (FDIC) regulated state bank. This practice allowed them to benefit from an established interpretation of the 1864 National Bank Act in which loans which cross state lines only had to obey the consumer protection laws of one state, those of the home state where the originating bank head office was located. Attempts have been made to pass a “Madden fix” act in Congress to correct this situation, but this has proved controversial—viewed by many members of Congress as a violation of state rights. The legal situation remains uncertain and the ability of platforms to lend into the 2nd circuit states of Vermont, Connecticut, and New York remains hampered.

3.2 Regulatory Structure and Borrower Protection

The US framework of financial regulation is by far the most complex in the world. This complexity arises is due to the interaction of various competing state and federal regulations and the presence of multiple regulators at the state and federal level.¹⁶ Alternative lending platforms are affected by a range of banking, securities, and consumer protection regulations at both state and federal level, regulations which continue to evolve through both new interpretations of case law and statute.¹⁷

For example, a long-standing characteristic of US banking, and a source of fierce debate over regulatory policy, is its so-called “dual banking system” (introduced under President Lincoln with the 1864 National Bank Act) with

¹⁵ *Madden v Midland Funding, LLC*, 786 F.3d 246 (2d Cir. 2015).

¹⁶ A long standing related discussion in law and regulatory scholarship concerns the degree to which the checks and balances of US state and federal regulation address the problems of “regulatory capture,” and how the various US regulatory agencies may best be organized and governed to support outcomes that are not excessively favoring particular interest groups (see [Baxter 2011](#)).

¹⁷ The discussion of the following paragraphs is largely based on ([Manbeck and Franson 2018](#)).

some banks chartered at state level, while other banks obtaining national charters. The state -chartered banks are regulated and supervised by state regulators. In addition, FDIC provides insurance to the banks that are not members of the Federal Reserve System. The national-chartered banks are supervised by the Office of Comptroller of the Currency OCC. The dual banking system also illustrates a central and also controversial aspect of US financial regulation, the practice of so-called Federal “preemption,” situations where Federal law allows financial services to be provided in multiple states under Federal law without having to comply with the financial regulations of each individual state.¹⁸ Preemption is though inevitably controversial representing as it does the outcome of a struggle for supremacy between state and federal regulation.

There is a similar division between federal and state securities regulation. Since the passage of the 1933 “Truth in Securities Act”, all securities sold and traded must be registered with the SEC. The only exemption is securities issued and sold only within an individual state and recommended by in-state investment advisers. Nationally registered securities though must still comply with all the state level securities regulation (so-called “blue-sky” laws) in those state where they are bought and sold, except for exemptions under certain specific rules which allow the sale to authorized (sophisticated) investors.

Similarly, alongside state level consumer protection laws and regulation, there is a federal level consumer protection framework operating through the Federal Trade Commission (FTC) and, since the 2010 Dodd-Frank act, the Consumer Financial Protection Bureau (CFPB). The Federal Trade Commission (FTC) act that “declares as unlawful any unfair or deceptive act or practice in or affecting commerce” (Manbeck and Franson 2018, section IVB). The FTC applies this to consumer lending through a range of credit practice rules that seek abusive lending. The CFPB further applies rules to enforce a range of other Federal legislation that seeks to prevent unfair, deceptive, and abusive acts and practices in consumer finance, including the Fair Credit Reporting Act, Truth in Lending Act, Equal Opportunity Act for credit reports, the Fair Debt Collection Practices Act. The Electronic Funds Transfer Act and Telephone Consumer Protection Act also impacts marketplace lenders.

Both the FTC and the CFPB are actively engaged in oversight of alternative lenders. The FTC recently brought a case against for LendingClub Corporation for falsely promising that it would not apply “hidden fees” to consumers (Federal Trade Commission 2018). CFPB since March 2016, has

¹⁸ See for example (Wilmarth Jr. 2004). An insightful panel discussion of dual banking and preemption is summarized by (Davis and Rice 2006).

started to accepted complaints on consumer loans from marketplace lending (Consumer Financial Protection Bureau 2016).

We now describe, some of the difficulties that have arisen in applying this regulatory framework to borrower protection. Unless a preemption applies, states require lenders to obtain licensing for brokering, lending, debt collecting, loan servicing, loan solicitation, and loan purchasing. Marketplace lenders have been able to avoid these state level requirements, by originating loans through a FDIC regulated or OCC national chartered bank and hence benefiting from the Federal preemption that requires compliance only with the consumer lending regulations of the bank's home state (that of its head office). For example both Prosper and LendingClub originate their lending through Webbank, an FDIC-regulated Utah bank.

There have been recent challenges at the state level which may limit this practice in the future. In 2017 and 2018, litigation upholding requirements to comply with state consumer protection regulation has been pursued in Colorado, California, Massachusetts and North Carolina. The key issue in these cases has been the wider applicability of the legal concept of "true lender" originally developed in the context of legal challenges to the operation of some payday lenders providing high-interest loans using a non-bank balance sheet lending business model. Is the "true lender" the marketplace lender or the FDIC-insured bank. Decisions from the courts have been inconsistent with some ruling that the originating bank is the true lender, whereas others concluding the marketplace lender is the true lender and therefore subject to all relevant state level laws and not benefiting from Federal preemption.¹⁹ The importance of this unresolved interpretation of regulation for US alternative lending platforms is illustrated by LendingClub highlighting in their 2017 10 K the potential damage their business model of an adverse interpretations of "true lending."

A related issue is the application of state level usury laws setting a maximum level of interest. The judgment that has had the most impact on US alternative lending platforms is the 2016 decision in the case of Madden

¹⁹ (Cadwalader 2018) provide insightful discussion, focusing on two prominent cases, one brought in Massachusetts against Kabbage Inc. the Alternative lending platform for small business loans and Celtic Bank Corporation, Kabbage's loan originator; another in Colorado against Avant and Webbank. They find that "In determining the 'true lender,' courts have developed two different approaches, with the choice of approach invariably dictating the result. Certain courts focus on the fact that the bank is the party to the loan agreement and is the entity that actually disbursed the proceeds. These courts conclude that the bank is the true lender and that federal preemption applies. Other courts, however, conduct a broader analysis, focusing on the origination and underwriting of a loan, as well as any material interest in the extension of credit, if any, the bank retains after origination. These courts conclude that the finance company is the true lender and, therefore, federal preemption does not apply." See also (Lo 2016b) for similar argument.

v Midland Funding LLC (this concerned a credit card debt that was first originated by Bank of America, but then eventually after a portfolio sale was then transferred to Midland Funding a non-bank vehicle used for management of defaulted debt).²⁰ The Second Circuit ruled that non-banks who received loans were not permitted to apply the same rate of interest that national banks can charge. Critics of this decision, concerned that this decision threatens the integrity of the US consumer credit and small business lending and loan-resale markets, argue that this is a misinterpretation because it does not recognize the long-established doctrine of “valid when made,” i.e., once a loan has been validly originated under applicable regulation then it can be resold without having to comply with further consumer protection regulations.²¹

3.3 Investor Protection

The initial years for US alternative lending platforms 2006–2008, in a deteriorating credit environment, were difficult for the nascent industry with high-default rates raising concerns about investor protection.²² The original contractual arrangement used then was similar to that still used by alternative lending platforms in the UK (see below), with investors directly holding either whole loans or shares in loans, and the platform playing the role of matching investors with borrowers and servicing the loans by collecting and allocating payments of interest and principal.

A key development in investor protection regulation was in November 2008, when the SEC entered a cease-and-desist order against Prosper stating that it had been selling securities which must under US securities law be registered with the SEC (see Lo 2016a).²³ In response, the legal compliance of platforms changed, with investors now required to hold non-recourse notes issued under a “shelf registration” with the SEC (i.e., a single registration covering multiple note issues), each note representing claims on the payments of borrower interest and principal for a set of specified loans.

²⁰ The Madden case has proved to be a useful natural experiment for economists investigating the impact of alternative lending on the supply of consumer credit, see (Honigsberg et al. 2017; Danisewicz and Elard 2018).

²¹ This perspective is forcefully argued by (Davis 2018), in a white paper commissioned by the MarketPlace Lending Association representing the major alternative lending platforms.

²² (Lo 2016a) reports from press coverage that “Even as the industry grew, lenders bore painfully high default rates—Prosper was charging off more than 20% of loans issued before 2008, while LendingClub fared better, but still had 8.5% of its pre-2008 loans in default.”

²³ This and follow paragraphs are again based substantially on (Manbeck and Franson 2018).

This was a material difference to their business models because. In order to sell these notes to a broad investor based, the platforms had to comply with the requirements for registration and provision of investor information under the Truth in Securities Act and Securities Exchange Act (1933) as well as complying with all requirements for offering advice in securities investment and further requirements if they support a secondary market for resale of investments. The substantial burden of SEC registration, together with the accompanying obligation to comply with investment advice and the “blue sky” securities laws of all states where investor notes were sold led to rapid industry consolidation. Consolidation left only two platforms surviving—those of Prosper and LendingClub.

(Manbeck and Franson 2018) highlight the compliance costs for marketplace platforms who offer notes, i.e., for alternative lending platforms when compared with balance sheet lenders.: “*The SEC registration process is complex, time-consuming and expensive. Operators who choose to register their Platform Notes for sale to the general public must be prepared to devote substantial resources to the effort.*”

Sale to the general public through SEC registered notes is not the only possible funding route for alternative lending platforms in the US. As both Prosper and LendingClub have demonstrated there is a substantial market for investment in platform loans among institutional investors. Sale of notes to institutional investors or other “accredited investors” does not require the compliance costs of shelf registration.²⁴

Meeting the “accredited investor” requirement is relatively easy for these larger platforms that can build or use existing investor relationships with institutional investors. However, it is difficult for smaller or startup platforms that are trying to establish themselves through accredited investor funding. Accredited investor funding enables the platforms to avoid the compliance costs of shelf registration, which can be prohibitively high at an early stage of platform development. An important further development was the amendment of rule 506 of Regulation D under the securities act in September 2013 which enabled the use of general advertising if notes are sold to “accredited investors.”

²⁴ As explained by (Manbeck and Franson 2018) “the term ‘accredited investor’ includes most institutional investors and individuals who (i) individually, or with their spouse, have a net worth exceeding \$1 million exclusive of the value of the person’s primary residence (and subject to certain adjustments for “underwater” mortgages), or (ii) individually had an income in excess of \$200,000 in each of the two preceding years, or had a joint income with spouse in excess of \$300,000 in each of those years, and have a reasonable expectation of reaching the same income level in the current year.”

A further compliance issue for US alternative lending platforms is their fiduciary duties as investment advisors which give rise to potential compliance requirements, under the Investment Advisor Act of 1940. They must show appropriate management of conflicts between their duties of providing best transaction execution and managing the client's portfolio in their best interest.²⁵ See Manbeck and Franson (2016).

One prominent case highlights the potential high costs of non-compliance with these fiduciary responsibilities. LendingClub Asset Management (LCAM, formerly known LendingClub Advisors LLC), is a subsidiary of LendingClub Corporation and a wholly owned investment advisor registered with the SEC. On September 28, 2018, the SEC (2018) announced the outcome of an investigation into fraud for using “*funds improperly using fund money to benefit LendingClub Corporation.*” SEC levied a combined penalties of \$4.2 million on LCAM and on Renauld Laplanche, the former CEO of Lending Club Corporation until his resignation in June 2016 over an unrelated breach of Lending Club business practices.

4 Regulation in China

4.1 Overview

China's first online lending platform, PPDAl group launched in 2007, and the industry has gone through rapid growth since then. It peaked at about 3,500 in number in 2015 and has become the biggest in the world valued at 1.3 trillion yuan (US\$190 billion) (Wildau and Jia 2018). In June 2018 alone, the industry channeled loans from 4.1 million investors to 4.3 million borrowers (Wildau and Jia 2018). China's P2P platforms usually set up and register themselves as consultancy firms and therefore the average size of them is small. As of the middle of 2014, the averaged registered capital of the firms was about RMB23.7 million and most platforms had registered capital between RMB 5 million to 20 million (Shen 2015). Only 54 platforms had registered capital of above RMB 50 million (Lufax 2015). As a result, the market is highly dispersed and competitive.

In contrast with the U.S, retail rather than institutional investors dominate China's P2P sector. For instance, in a survey by ACCA (2015, pg 17), most investor respondents said that they were carrying out investments on their own behalf (85%). A further 3% said they were also investing for family and

²⁵ See.

friends, while only 1% of lenders said they were investing for others on an institutional basis.

Judging by the geographic distribution of the platforms, there is a direct correlation between the number of local platforms, local economic development and private lending activity. The P2P platforms are concentrated in four economically developed regions, namely, Guangdong, Zhejiang, Beijing, and Shanghai, which together account for 63 percent of all platforms in China (China's Internet Lending Business Annual Report 2016). Households and SMEs have been the principal borrowers. Around 80% of P2P loans are less than RMB100,000, with SMEs accounting between 20 to 40% of all loans (ACCA 2015).

A major factor accounting for the explosive growth is that from 2007 to July 2015 the online lending market was largely left unregulated. This laissez-faire approach attracted millions of investors rushing in the online lending business and a fertile soil for enormous irregularities and fraud and increasing numbers of platform failures.²⁶ In July 2015 the People's Bank Of China (The PBOC) issued their plans for regulatory reform (People's Bank of China 2015) referred to as the "Guiding Opinions." Since then, in an still ongoing process, the Chinese authorities have been developing a comprehensive, systematic regulatory regime governing the P2P sector.

The new regulatory framework creates a dual supervision module under which the regulatory responsibilities are shared between the China Banking and Insurance Regulatory Commission (CBIRC) and provisional governments. It imposes three key requirements on all remaining platforms: a requirement of recordation and registration; a requirement of custodian; and a requirement of mandatory information disclosure. These requirements considerably raised the threshold of online lending market and will inevitably lead to a substantial restructuring of the sector.

4.2 2007–2015: No Regulation

In the absence of regulation, why were Chinese investors still willing to take the risk of lending their money to complete strangers? The answer lies in the unique features of China's P2P business models. In China, in order to attract investors, most platforms offered investors "guarantees," often of both

²⁶ Since 2011, increasing number of P2P lending platforms have abruptly collapsed due to financial difficulties, leading to the CBIRC issued "Circular on Risks Associated with Peer-to-Peer Lending" on August 23, 2011. However, instead of providing specific rules on P2P lending, the Circular merely identifies a number of risks associated with P2P lending as a result of lack of regulation and supervision, such as illegal funding, fraudulent activities.

principal and interest, backed by the platform itself, a financial institution, or a guarantee company about which the investors know little to nothing (P2P Research Group 2016, pg 8). The platforms regularly and invisibly paid out their own funds to reimburse investors for any losses, making defaults “disappear” (P2P Research Group 2016, pg 8).

Another key common feature China’s P2P business models was that the platforms often originated and held loans. As a result, there was no direct contractual relationship between the lenders and borrowers, instead, they transacted with the platform separately (P2P Research Group 2016, pg 9). China’s P2P platforms were acting as banks but without any banking regulation.

This raised severe issues of investor protection. Investors were seriously misled by the “guarantee” and failed to understand the risk associated accompanying P2P lending, such as platform collapse and borrowers default. On the other hand, there was no deposit insurance for P2P investors, nor were there capital reserve requirements or any other equivalent rules with which P2P lending platforms had to comply. As revealed later in the crisis, many platforms were actually scams from the beginning. Borrowers also suffered from unfair lending and collection practice, such as potentially unclear or misleading lending terms, predatory or discriminatory credit decisions, and abusive or deceptive serving acts (Shen 2016, pg 205). The Chinese government was well aware of the potential risks that were likely to arise from the China P2P sector (Shen 2016), but it chose to turn a blind eye to it till July 2015, following the emergence of hundreds of platform scandals.

4.3 The New Regulatory Framework

The Chinese government’s hands-off approach to P2P lending industry has dramatically shifted since July 2015. A comprehensive, systematic regulatory regime governing the P2P sector now has been established on basis of four key documents (referred to “1 + 3” framework). These are their Interim Measures on Online Lending (China Banking Regulatory Commission 2016a), Guideline on Online Lending Registration (China Banking Regulatory Commission 2016b), Guideline on Custodian Business (China Banking Regulatory Commission 2017a), and Guideline on Information Disclosure (China Banking Regulatory Commission 2017b).

Of these the 2016 Interim Measures is the corner stone of the regulatory regime. It contains a total of 47 articles that cover all important aspects of the online lending industry. Contravention of the 2016 Interim Measures can lead to administrative or even criminal penalty. The three guidelines are

implementing rules of the 2016 Interim Measures, each focuses on a particular issue. Except for these regulations that specific to online lending industry, P2P platforms are also subject to general legal rules, in particular Company Law, Contract Law, Anti-competition Law, Consumer Protection Law and Anti-Money Laundering (AML) Law.

4.3.1 The Regulatory Model

The regulatory model created by the 2016 Interim Measures is often described as a “dual supervision” model. At central level, the current financial regulatory framework in China is sector-based, with sector-specific regulators responsible for banking and insurance and securities, respectively: the China Banking and Insurance Regulatory Commission (CBIRC)²⁷; and the China Securities Regulatory Commission (CSRC). The 2016 Interim Measures reaffirmed that online lending falls within the jurisdiction of the CBIRC.²⁸ According to Article 4, the CBIRC acts as the lead regulator for China’s P2P industry, empowered to develop rules for supervision and administration of the business activities of platforms as well as carry out regulation of business conduct. In the meanwhile, provincial governments²⁹ are authorized to supervise P2P platforms within their respective jurisdictions, in particular regarding their recordation and registration.³⁰ They also have the power to impose penalties on misbehaving platforms.³¹ Any particular platform is therefore subject to a dual supervision of the CBIRC and provincial government where it is registered.

In addition, three other government agencies also have supervisory power over P2P lending platforms in relation to particular issues.³² The Ministry of Industry and Information Technology is responsible for conducting supervision over the telecommunications business involved in online lending business activities. The security of internet services provided by online lending platforms is supervised by the Ministry of Public Security.³³ The

²⁷ Until April 8, 2018, banking and insurance industry was regulated by the China Banking Regulatory Commission (CBRC) and the China Insurance Regulatory Commission (CIRC) respectively. These two regulatory bodies were then merged to form the CBIRC.

²⁸ The division of regulatory responsibilities was first made by the 2015 Guiding Opinions.

²⁹ More specifically, this responsibility is assumed by an agency called Financial Work Office.

³⁰ Article 4 and Article 5.

³¹ Article 40.

³² Article 40.

³³ Its power includes investigate and punish violations, and crack down on financial crimes involved in online lending.

State Internet Information Office has the power to conduct supervision over financial information services and internet information contents.

The current regulatory framework also has an element of self-regulation. The National Internet Finance Association (NIFA) run by the central bank has 400 members recruited to regulate China's P2P sector. Article 34 of the 2016 Interim Measures delegates the NIFA the power to establish rules of self-regulation and industry standards, to accept complaints and reporting, to conduct self-inspection, and so on. Given the limited resources of regulators and the challenging nature of the P2P industry, the NIFA is expected to play an important role under the current regime.

4.3.2 Registration of Platforms

Previously, China's P2P platforms usually registered as consultancy firms and therefore only subject to the general rules of Company Law and Regulation of Company Registry. Article 5 of the 2016 Interim Measures introduced a new procedure to set up P2P platforms. First, it needs to get a business license from the company registry; second, it needs to conduct recordation and registration with the local financial regulatory authority at the place where it is based; third, apply for a relevant telecommunications business permit from the competent communication agency. The new procedure not only applies to new platforms, also applies to existing platforms. An existing platform will be prohibited from continued operation if it fails to go through the new procedure. The first and third steps are general in nature and not difficult to get, the second step, however, is very challenging to all existing platforms.

Chapter "[Innovations in Alternative Finance in Historical Perspective](#)" of the 2016 Guideline on Online Lending Registration details the supporting documents and a timeline. On August 13, 2016, the Task Force of Online Lending Rectification issued a "Checklist for the Compliance of Online Lending Information Intermediaries" (108 Checklist) which sets up a uniform standard for recordation and registration for all P2P platforms in China and a new deadline of December 2018.³⁴

The 108 Checklist aims to clean up the irregularities and illegalities prevailing in P2P sector and to restore the identity of platforms as real online information intermediaries. Past practices like pooling funds, guaranteeing investors' return, issuing wealth management products will no longer be tolerated by the regulator. The majority of existing platforms are unlikely to complete the registration and therefore will be forced out of the market. The

³⁴ By January 2019 no single platform had completed the procedure.

estimation of surviving numbers ranges from 50 (Weinland and Ju 2018)³⁵ to 400 (Cao 2018).

4.3.3 Borrowers Protection and Responsible Lending

The current regulation devotes very little attention to borrowers. Although Chapter “[An Introduction to Cryptocurrencies](#)” of the 2016 Interim Measures is entitled with “Protection of Lenders and Borrowers,” borrowers protection is only briefly mentioned in Article 27 which provides that platforms must ensure the information of borrowers are safely stored and properly used. The protection of borrowers thus is largely left to other laws and regulations. Among which, the PBOB’s 2002 Circular on Closing Down Underground Banks and Combating Usury (the PBOB’s 2002 Circular) and the Provisions of Supreme People’s Courts on Certain Issues Concerning Application of Law in Trial of Cases involving Private Lending³⁶ (The Supreme Courts’ Provisions on Private Lending) are the two major documents that governing P2P lending.

The PBOC’s 2002 Circular provides that interest rate in private lending is decided between borrowers and lenders, up to four times the interest rate charged on similar types of bank loans during the same period.³⁷ Since P2P lending is regarded as private lending in China, The PBOC’s 2002 Circular applies. But this four times limit was relaxed by the Supreme Courts’ Provisions on Private Lending in 2015. The Supreme Court provides that

- i. An annual interest rate not exceeding 24% will be held legally binding and enforceable by a court
- ii. An annual interest rate exceeding 36% will be considered to be illegal
- iii. Any interest paid by the borrowers and exceeding the annual interest rate of 36% shall be refunded to the borrower
- iv. While an agreement on an annual interest rate between 24 and 36% constitute a “natural obligation” under Chinese law, which means both (1) the lenders’ request of the payment of the interest exceeding the interest rate of 24%, and (2) the borrower’s request of the refund of any already paid interest exceeding the interest rate of 24% will not be upheld by a court.

³⁵ According to Mr Greg Gibb, Chief executive of Lufax, see (Weinland and Ju 2018).

³⁶ Issued by the Supreme Court on August 6, 2015.

³⁷ Above this four times limit will be regarded as usury.

4.3.4 Retail Investor Protection

Since China's P2P lending industry is dominated by retail investors who suffered substantial losses in recent crisis, investor protection has been a central objective of the new regulations. The key elements include imposing lending limits and custodian requirement and strengthening information disclosure.

- **Lending Limits** P2P platforms are requested to follow the principle of “diversification and small amount” in loan making. Article 17 of 2016 Interim Measures now provides that the balance of loans of the same natural person on one online lending platform cannot exceed RMB 200,000; for a legal person or any other organization, the upper limit is set as RMB 1 million. There are further caps on the total balance of loans obtained by the same person from different online lending platforms: for a natural person, it is capped at RMB 1 million, for a legal person it is capped at RMB 5 million.
- **Custodian Requirement** As a response to the scandals in which many platform owners absconded with funds, Article 28 of 2016 Interim Measures provides that an online lending platform must separate its own funds from the funds of lenders and borrowers. The funds of lenders and borrowers must be in custody of a qualified banking financial institution. In 2017 Guideline on the Custodian Business, it further states that fund custodian services can only be provided by commercial banks³⁸; a commercial bank should set up special custodian accounts and cannot outsource the relevant work like account opening, trading password verification³⁹; one online lending platform can only be served by one custodian.⁴⁰ This is a further incentive for consolidation since many small platforms will have difficulty in finding a custodian.
- **Information Disclosure** In the past, there was no mandatory requirement or standard form for information disclosure. The 2016 Interim Measures devotes the whole Chapter “[The Role of Utopia in the Workings of Local and Cryptocurrencies](#)” to the issue of information disclosure and 2017 Guideline on Information Disclosure provides further details. A platform is now under an obligation to make truthful, adequate, complete and timely

³⁸ Article 5.

³⁹ Article 12.

⁴⁰ Article 14.

information disclosure to the public in relation to the platforms, its operation, loans, other material risks and channels for consumer complaint.⁴¹ A platform is obligated to disclose basic information regarding the borrowers and risk assessment prior to investment and to make continual disclosure regarding borrowers' financial status and factors that may affect their repayment capability.⁴² These information disclosure announcements and relevant documents are submitted periodically to the local government where it is registered and placed in its registry house for public inspection.⁴³ In addition, third party intermediaries like accounting firms, law firms and information security firms are recruited to ensure the quality of information.⁴⁴

- **Fraud** Fraud is the most serious concern for investor protection and for the survival of China's P2P sector. The 2016 Blue Book of Internet Finance found 896 P2P platforms that got into troubles in 2015, with more than half involved in fraud. Many platforms were set up as scams from the outset. In the most notorious case of Ezubao, the platform was found to fabricate 95 percent of its projects (Caijing 2018).⁴⁵ The recent regulations of imposing lending limits, funds custody and information disclosure, are all designed to prevent fraud. Under the current Criminal law, fraudsters can be charged with the crime of fraud on fund raising,⁴⁶ or illegal pooling of public deposits,⁴⁷ both can lead to a death penalty.
- **Fiduciary Duties** According to Article 148 of the 2006 Chinese Company Law, member of the board of directors, member of the board of supervision, senior managers are all owed a duty of loyalty and duty of care to their company. Article 150 further states that directors, supervisors and senior managers are liable for the losses resulted from their breach of the law, administrative regulation or article of association. The same rules apply to P2P platforms as well. Unfortunately, the 2006 Company Law is notoriously vague on fiduciary duties. It fails to spell any details, above all, the standard of fiduciary duties. As a consequence, fiduciary duties are rarely enforced in practice. This is even more so for P2P platforms. As most of the problematic platforms are involved in some criminal offences, investors

⁴¹ 2017 Guideline on Information Disclosure, Article 2; Chapter "Innovations in Alternative Finance in Historical Perspective" specifies the content of the required information.

⁴² 2017 Guideline, Article 9.

⁴³ 2016 Interim Measures, Article 31.

⁴⁴ 2016 Interim Measures, Article 31.

⁴⁵ Ezubao was once one of the biggest platforms, by December 8, 2015, its trading volume totaled at 74.568 billion yuan, involved investors 909,500, see (Caijing 2018).

⁴⁶ Criminal Law, Article 192.

⁴⁷ Article 176.

normally go to the police as opposed to the court seeking for redress. Given the weakness of the law and the dispersed nature of China's P2P investors, fiduciary duties are not very likely to be enforced at present or in the near future.

5 Other Countries

This section offers a brief review of the law and regulation of alternative lending in other countries.⁴⁸ For a more detailed discussion we refer to the various sources we have consulted.⁴⁹

There is substantial variation across countries in the regulatory regimes for alternative lending. Only a few have an explicit regulatory framework. One country which does is France where a 2014 crowdfunding regulation created two categories of platforms: the CIP or *conseil en investissement participatif* which plays an advisory role and has been used for equity-based crowdfunding and the IFP or *intermédiaire en financement participatif* which is a more passive and used for loan-based crowdfunding. Lending through IFP offers an exemption to the normal requirement for a banking license in order to provide small business loans (but this does not apply to consumer lending which still requires a banking license).

Other examples of tailored regulatory regimes are New Zealand, which introduced a regulatory framework for marketplace lenders starting in 2014, supporting strong growth in alternative consumer lending platforms; the Netherlands where a 2016 crowdfunding framework was introduced, supporting both equity and loan-based crowdfunding and allowing “dispensation” from securities laws for approved platforms leading to relatively strong growth especially of alternative lending to small business; and Spain where a 2015 act has led to some 11 alternative lending platforms being registered with the Comisión Nacional del Mercado de Valores (CNMV) the national securities regulator.

⁴⁸ While we have looked at Europe, some countries and Asia and North America, we have made no effort to review the regulatory framework in Latin America.

⁴⁹ Three sources provide convenient comparative reviews of the law and regulation of crowdfunding in different countries: (European Crowdfunding Network 2017) for Europe, Northern America and Israel; (European Commission 2017) for the 28 EU countries and for responses to an OECD survey on regulation of loan based crowdfunding (Havrylchuk 2018). In preparing this section we have also consulted (Havrylchuk 2018; Ahern 2018; Pranjivan 2017; Klöhn 2018; Ferrarini 2017; Milne 2018).

Most other jurisdictions have, as yet, have no explicit regulatory framework. This reflects a variety of circumstances—sometimes a conscious decision to support alternative lending using the existing framework with a cautious approach that limits development or simply a “regulatory lag” (Ahern 2018), i.e., a failure of the regulation to keep pace with market innovations.

A prominent example of conscious adaption of the existing regulatory framework in Australia, where the existing regulatory framework for “managed investment schemes” has been extended to cover a number of newly established alternative lending platforms offering investment opportunities to professional and sophisticated investors.⁵⁰ This brings these platforms under the jurisdiction of Australian Securities and Investments Commission (ASIC) the securities regulator rather than the Australian Prudential Regulation Authority (APRA) the banking and insurance regulator. The regulatory framework pays close attention to investor protection. While it allows platforms to apply for permission to offer investment opportunities to retail investors, approval is far from automatic forthcoming only when ASIC is persuaded that these investment opportunities are appropriately tailored and sufficiently clearly presented to be suitable as retail investment products. As of late 2017 only two retail platform approvals had been given. Canada has pursued a similarly cautious approach to Australia, supporting institutional investment in alternative lending platforms but to date authorizing only one alternative consumer lender available to retail investors (Hutchison 2018).

In some jurisdictions (examples include Belgium, Greece, Hungary, Italy, Luxembourg, Malta, Portugal) where lending to consumers or small business requires a banking licence, there appears to have been virtually no development of alternative non-bank lending of any kind and no lobbying of regulators or government to allow create a framework to support alternative lending platforms.

In other jurisdictions which similarly require a banking license to provide the consumer of small business loans, alternative lending platforms have emerged. This has happened in different ways. In Austria, Germany, Slovenia, Slovakia and Poland regulation allows the offer of subordinated loans to business without a banking license and without the need to be fully compliant with securities law such as the requirements of the EU prospective directive. As a result, a number of alternative lending platforms have emerged supporting investment in subordinated loans.

⁵⁰ For more detailed review see (Milne 2018).

Among these countries which require banking licenses for consumer and small business lending, only Germany has active alternative consumer lending platforms. These meet the requirement for a banking license, by following a similar business model to US platforms: loans are first originated by a bank and then transferred onto a platform for investor funding. Possibly other jurisdictions will follow suit, allowing the emergence of alternative consumer lending platforms, but this may require changes in law or regulation.

The development of alternative lending platforms is easier in more permissive jurisdictions (examples include the Czech Republic, Ireland, Finland, Gibraltar, Malta, Switzerland) where non-bank lending is already possible. Broadly, in these jurisdictions consumer lending requires only a non-bank consumer credit license and business lending requires only a general authorization as a financial services firm. Switzerland has pursued a particularly supportive approach, allowing platforms exemptions from securities and payments regulations with the goal of promoting small business lending. In a number of other jurisdictions—where alternative lending is in its infancy—the regulatory situation remains unclear.

The European Commission has recently proposed a crowdfunding directive covering both equity and loan-based crowdfunding (The European Commission 2018). The scope of this directive is limited to platforms that support equity investment, loans, and other forms of transferable investment in small- and medium-sized enterprise; it does not cover platforms supporting unsecured consumer borrowing. It defines investors as “persons” and as proposed appears to discourage institutional investment (“This Regulation aims to facilitate direct investment and to avoid creating regulatory arbitrage opportunities for financial intermediaries regulated under other Union legislation, in particular Union rules governing asset managers. The use of legal structures, including special purpose vehicles, to interpose between the crowdfunding project and investors, should therefore be strictly regulated and permitted only where it is justified.”) Crowdfunding platforms will be registered with and supervised by The European Securities Market Authority ESMA. In place of the usual prospectus requirements, platforms are required to produce a key investment information sheet, a maximum of 6 pages when printed on A4 paper without footnotes and with a mandatory statement emphasizing that the investments are high risk not savings products and not covered by any form of deposit insurance.

In the light of the review of law and regulation of alternative lending provided in our chapter, the proposed regulation seems unlikely to do much to promote alternative lending in the European Union. It is orientated primarily to equity-based crowdfunding, requiring for example platforms

to ascertain investor's previous experience with crowdfunding investments. There is little further attention to the key issue of retail investor protection, beyond the requirement for the investment information sheet.

Turning to Asia (outside of Australia and New Zealand which have already been discussed), Japan has a 2014 law providing a regulatory framework for equity crowdfunding, but no corresponding framework for alternative lending. A platform originates the loan with funds it collects from investors through silent partnership contracts (*tokumei kumiai keiyaku*: or TK agreements) under the Commercial Code of Japan as a business operator and the investors then provide funds to the P2P lending platform as silent partners (Samitsu 2017). This arrangement though has the disadvantage that, in the event of platform bankruptcy, investors would end up with a general claim on the platform there is no arrangement for continued separate servicing of their loan investments. Korea, in contrast, has allowed alternative (P2P) lending to take place with little supervision or regulation. Concern about rising default rates and lack of investor protections has resulted in an ongoing review by the Financial Services Commission and the likely introduction of a much stricter framework (Eur-Young 2018).

Indonesia introduced P2P lending regulation in 2017 (Chang 2018). This appears to have been a response to the establishment of a relatively large number of platforms under what was previously a regime of limited regulation, heeding some of lessons from the growing problems in Chinese P2P lending. In this regulation P2P lending is defined as "providing financial services to match a lender with a borrower to enter into a loan agreement in Rupiah currency directly through electronic systems using internet networks." Platforms must register with the Otoritas Jasa Keuangan (OJK) the financial regulator and subject to supervision to promote investor protection.

P2P lending started in India around 2012. India had no specific laws for P2P until 2017 when the Reserve Bank of India (RBI) issued its P2P lending Platform Directions (2017), coming into force that October 2017 (Nahata 2018). These require registration with the RBI subject to strict operational criteria on the scope of activities, prudential norms, operational guidelines, fair practice code, and reporting requirements. The directions cap limits lenders' exposure to Rs 10 lakh potentially deterring institutional investing. On the other hand, P2P lenders using a registered Non-Bank Financial Company will be able to obtain credit bureau data to access credit scores to evaluate risk. Reportedly, following the introduction of the directions, the number of alternative lending platforms in India has fallen from over forty to only five (Nahata 2018).

What overall conclusions can be drawn? Our brief review indicates the extremely wide range of law and regulations applied to alternative lending platforms across different jurisdictions. In some, e.g., France, regulation is orientated very much to the support of alternative business lending platforms. Others support a combination of consumer and business lending. All jurisdictions though need to address much the same concerns as those we document in the US, in China, and in the UK. For consumer loan origination appropriate consumer credit regulation is still needed, whether this is through requiring origination by a licensed bank or offering appropriate non-bank consumer credit license. Investor protection is also everywhere a central concern. Where regulation does not already impose strict requirements to ensure that alternative lending investments are not sold inappropriately to uninformed retail investors (e.g., Australia), re-regulation to protect investors can be expected.

6 Conclusions and Recommendations

The rise and continuing rapid growth of alternative lending platforms—in the US, China, and the UK and more recently in other countries—has been heralded as a fundamental technological challenge to traditional bank business models. Some have envisaged the emergence of direct platform-based lending as the early stages of widespread disintermediation, supplanting bank-intermediated lending to households and small businesses with direct funding of loans by personal savers offered more attractive short-term investment opportunity than that provided by regulated bank savings deposits.

Our review of law and regulation of alternative lending platforms in this chapter suggests a more cautious view of the development of this new industry. The rapid growth of the original and much-heralded form of “P2P” platform model with investment funded directly by households, appears more as the exploitation of a regulatory gap than as a technologically driven new business model. In most countries, until very recently there has been no coherent investor protection framework covering the potentially risky direct participations in household and small business loans. These are investments that can potentially prove unsuitable for unsophisticated investors.

Regulation and regulators have been slow to catch up with the new business models of alternative lending. How this happens is and will be critical for the future development of alternative lending. Our documentation of the experience of the US, China, and the UK show that regulation when eventually getting to grips with alternative lending has led more to change in the

business models of alternative lending platforms than in the business models of banks.

The US was the first jurisdiction to respond to potential investor risks in alternative lending, with its 2008 decision to treat platform investments as securities that must be registered with the SEC and subsequent development of federal regulations for investor protection in marketplace lending. These developments led both to the closure of many of the first US platforms and to a shift in the two major platforms that remained—Prosper and Lending-Club—from retail to institutional investor base which now accounts for more than 80 per cent of their loan funding.

A similar tightening of investor protection and other regulations is now taking place in the UK and in China. In China much of the so-called “P2P” lending has not in fact been alternative platform-based lending at all. There has been comingling of investor and platform assets and undeliverable return guarantees, similar to those offered on regulated bank deposits; but without the support of appropriate controls of credit risks, of any deposit insurance, or anything like sufficient equity capital to absorb losses, i.e., platforms have operated as poorly run banks not as alternative lending platforms at all. As the Chinese regulators get a grip on this situation, a substantial reintermediation is taking place with a sharp contraction in both the number of platforms and in the volumes of direct household to borrower lending they facilitate.

In the UK the growth of alternative lending platforms has also benefited from a relatively “light touch” regulatory regime, one that until 2016 received strong political support at the highest levels of government. This situation is though now changing, with growing concerns about investor risks leading to the FCA review and consultation on new regulations set out in (FCA 2018) and proposed for introduction in 2019 proposing substantially closer oversight of platforms and business models. In the UK as in the US, regulation is pushing alternative lending toward business models that are substantially funded by institutional investors who can fully assess the underlying risks rather than by individual retail investors.

The reregulation of alternative lending platforms seems to us to be an unavoidable growing pain, a response to the development of their new business models that will inevitably take place eventually also in other countries. As the industry develops the framework of regulatory oversight has to respond, taking account of the limited ability of the majority of retail investors to assess and understand the risks of this investment class and hence introducing stronger retail investor protections. A shift toward an institutional rather than retail funding base is a natural and inevitable further consequence of this reregulation.

This also suggests that the efforts by the European Commission to promote crowdfunding, through the new proposed crowdfunding regulation (The European Commission 2018), will do little or nothing to improve the supply of credit to small and medium-sized businesses in Europe (one of the supposed motivations for this regulation). The perspective we have developed challenges the Commission's approach in several respects: (i) they conflate the very different activities of equity and loan-based crowdfunding when these need to be treated differently; (ii) it has missed an opportunity to support the supply of short-term credit to small and medium-sized business through the promotion of the broad category of marketplace lending, both from non-bank balance sheet lenders and from alternative lending platforms, restricting attention to the very narrow and less obviously viable sub-division of peer-to-peer (crowdfunded) lending by individuals to businesses; (iii) It has failed to recognize the need to support the most natural source of funding of alternative lending platforms operate at pan-EU level, that provided by investment funds and other financial institutions, a major omissions since it will be relatively difficult to effectively assess and communicated risks on small business lending to retail investors on a pan-EU basis limiting the potential for retail funding.

We conclude on one final point. Must reregulation to promote retail investor protection prevent alternative lending ever becoming a major challenge to bank intermediation? We anticipate that this will be the case. But there is a technology-based opportunity that could lead to a different outcome. We observe that this re-regulation has been largely based on a traditional, manual rather than an automated approach to the regulation of a new and still small fintech-based sector. Alternative lending platforms are becoming subject to much the same burden of compliance on customer and investor protection as traditional banks or investment funds, with extensive reporting requirements and supervision and the eventual sanction of the removal of authorization in the event of non-compliance.

Going forward the prospect seems to be that they will be tightly regulated activities with much the same high costs of regulatory compliance faced by established financial services firms. While, unlike banks, they do not have to comply with prudential regulation requirements, on capital, liquidity, ring-fencing, etc., they do not enjoy the regulatory benefit of the deposit insurance arrangements. This is undermining much of the competitive advantage enjoyed by alternative lending platforms and hence limit the potential for technologically driven benefits to bank customers through disruption of traditional bank business models.

While there is an understandable imperative of protecting retail investors, it is striking that regulators seem to be making little or no attempt to take advantage of the new emerging financial technologies for regulatory compliance (“RegTech”) when addressing the regulation of alternative lending platforms, one of the most prominent of the new financial technology-based activities that are promoting welcome competition in bank lending and personal investment. The alternative approach to regulation that we believe should be explored further, is based on using the same emerging technologies that support borrowing and investing on alternative lending platforms in the regulation of both bank and non-bank lending. The key, we argue is ensuring that loans and payments of interest, capital, fees, and where applicable taxation are recorded on a completely granular and standardized basis, including a full range of accepted standard metrics of loan quality and performance, with full transparency to outsiders whether investors, advisers, regulators, or other interested parties limited only by the requirements of borrower confidentiality.

A granular data approach can allow third parties—for example, the money advice pages in newspapers and online media, also investment advisers—to compare the risks and returns of portfolios on different platforms. Provided the underlying loan portfolios are relatively low risk, then this would allow platform investment opportunities and risk to be adequately summarized and communicated by investment advisers and by financial journalists. Third parties could provide an evidence-based comparison of prospective investment performance on different platforms and the different risk classification they offer and this would allow media commentators to in turn explain investment opportunities in simple plain English.

References

- ACCA. 2015. *The rise of peer-to-peer lending in China: An overview and survey case study*. https://www.accaglobal.com/content/dam/ACCA_Global/Technical/manage/ea-china-p2p-lending.pdf.
- Ahern, Deirdre M. 2018. *Regulatory arbitrage in a fintech world: Devising an optimal EU regulatory response to crowdlending*.
- Baxter, Lawrence G. 2011. Capture in financial regulation: Can we channel It toward the common good. *Cornell Journal of Law and Public Policy* 21. HeinOnline: 175.
- Cadwalader. 2018. *Marketplace lending update: Who’s my lender?* Cadwalader, Wickersham and Taft LLP. <https://www.cadwalader.com/uploads/cfmemos/9cf7d89d0e9c1e0ddaebc58b8e8ecc3.pdf>.

- Caijing. 2018. *The dismal end of Ezubao*. May 28, 2018. <http://finance.caijing.com.cn/20180528/4459893.shtml>.
- Cao, Pei. 2018. Deadline for recordation and registration comes close, the last one kilometer for 10 years' P2P industry. *Phoenix Weekly*, April 5, 2018. <http://www.ifengweekly.com/detil.php?id=5612>.
- Chang, Soonpeel Edgar. 2018. *Regulation of crowdfunding in Indonesia*. Growth, 5.
- Chen, D., A.S. Kavuri, and A. Milne. 2020. Lessons from the rise and fall of Chinese peer-to-peer lending. *Journal of Banking Regulation*. <https://doi.org/10.1057/s41261-020-00132-2>.
- China's Internet Lending Business Annual Report. 2016. osscdn.wdzj.com/pdf/20170106.pdf.
- China Banking Regulatory Commission. 2016a. *Interim measures on online lending*. The ministry of industry and information technology, The ministry of public security, The State Information Office. http://www.cbrc.gov.cn/govView_37D312933F1A4CECBC18F9A96293F450.html.
- China Banking Regulatory Commission. 2016b. *Guideline on the administration of recordation and registration of online lending information intermediary institutions*. http://www.cbrc.gov.cn/govView_E7B94B41E8C340E4833472632308AEC5.html.
- China Banking Regulatory Commission. 2017a. *Guideline on custodian business for online lending funds*. http://www.cbrc.gov.cn/govView_4201EF03472544038242EED1878597CB.html.
- China Banking Regulatory Commission. 2017b. *Guideline on information disclosure by online lending information intermediary institutions*. http://www.cbrc.gov.cn/govView_C8D68D4C980A4410B9F4E21BA593B4F2.html.
- Consumer Financial Protection Bureau. 2016. *CFPB now accepting complaints on consumer loans from online marketplace lender*. <https://www.consumerfinance.gov/about-us/newsroom/cfpb-now-accepting-complaints-on-consumer-loans-from-online-marketplace-lender/>.
- Danisewicz, Piotr, and Ilaf Elard. 2018. *The real effects of financial technology: Marketplace lending and personal bankruptcy*. Available at SSRN 3208908.
- Davidson, Jonathan. 2016. *Balancing regulatory objectives in the dynamic consumer credit market*. London: Financial Conduct Authority. <https://www.fca.org.uk/print/news/speeches/balancing-regulatory-objectives-dynamic-consumer-credit-market>.
- Davis, Erin, and Tara Rice. 2006. *Federal preemption of state bank regulation: A conference panel summary*. Chicago Fed Letter, no. 230a. Federal Reserve Bank of Chicago: 1.
- Davis, Polk. 2018. *Federal banking regulators can and should resolve madden and true lender developments*. New York. <https://www.davispolk.com/publications/federal-banking-regulators-can-and-should-resolve-madden-and-true-lender-developments>.

- Eur-Young, Hwang. 2018. FSC reins in P2P lending. *Korea JoongAng Daily*, June 18, 2018. <http://koreajoongangdaily.joins.com/news/article/article.aspx?aid=3049440>.
- European Commission. 2017. *Identifying market and regulatory obstacles to cross-border development of crowdfunding in the EU*. Brussels. https://ec.europa.eu/info/sites/info/files/171216-crowdfunding-report_en.pdf.
- European Commission. 2018. *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on European crowdfunding service providers (ECSP) for business*. <https://ec.europa.eu/transparency/regdoc/rep/1/2018/EN/COM-2018-113-F1-EN-MAIN-PART-1.PDF>.
- European Crowdfunding Network. 2017. *Review of crowdfunding regulation*. Brussels. https://eurocrowd.org/wp-content/blogs.dir/sites/85/2017/10/ECN_Review_of_Crowdfunding_Regulation_2017.pdf.
- FCA. 2016a. *Call for input to the post-implementation review of the FCA's crowdfunding rules*. London. <https://www.fca.org.uk/publication/call-for-input/call-input-crowdfunding-rules.pdf>.
- FCA. 2016b. *Guide for consumer credit firms*. <https://www.fca.org.uk/publication/financialised-guidance/consumer-credit-being-regulated-guide.pdf>.
- FCA. 2016c. *Interim feedback to the call for Input to the post-implementation review of the FCA's crowdfunding rules*. 16–13. Feedback Statement. London. <https://www.fca.org.uk/publication/feedback/fs16-13.pdf>.
- FCA. 2018. *Loan-based ('peer-to-peer') and investment-based crowdfunding platforms: Feedback on our post-implementation review and proposed changes to the regulatory framework*. London. <https://www.fca.org.uk/publications/consultation-papers/cp18-20-loan-based-peer-peer-and-investment-based-crowdfunding-platforms-feedback-our-post>.
- FCA. 2019. *Loan-based ('peer-to-peer') and investment-based crowdfunding platforms: Feedback to CP18/20 and final rules*. London. <https://www.fca.org.uk/publications/policy-statements/ps19-14-loan-based-peer-to-peer-investment-based-crowdfunding-platforms-feedback-final-rules>.
- Federal Trade Commission. 2018. *FTC charges lending club with deceiving consumers*. <https://www.ftc.gov/news-events/press-releases/2018/04/ftc-charges-lending-club-deceiving-consumers-0>.
- Ferrarini, Guido. 2017. Regulating fintech: Crowdfunding and beyond. *European Economy* (2). Europeye srl: 121–142.
- Havrylychuk, Olena. 2018. *Regulatory framework for the loan-based crowdfunding platforms*. OECD.
- HM Revenues & Customs. 2015. Income tax: Innovative finance individual savings account and peer to peer loans—GOV.UK. [Www.Gov.Uk. https://www.gov.uk/government/publications/income-tax-innovative-finance-individual-savings-account-and-peer-to-peer-loans](http://www.gov.uk/government/publications/income-tax-innovative-finance-individual-savings-account-and-peer-to-peer-loans).
- Honigsberg, Colleen, Robert J. Jackson Jr., and Richard Squire. 2017. How does legal enforceability affect consumer lending? Evidence from a natural experiment. *The Journal of Law and Economics* 60 (4): 673–712.

- Hutchison, Connor. 2018. The future of peer-to-peer lending in Canada. *Ivey FinTech Club*. <https://medium.com/iveyfintechclub/the-future-of-peer-to-peer-lending-in-canada-66db882f5835>.
- Klöhn, Lars. 2018. The regulation of crowdfunding in Europe. In *The economics of crowdfunding*, ed. Douglas J Cumming and Lars Hornuf, 219. Springer.
- Lo, Benjamin. 2016a. It ain't broke: The case for continued SEC regulation of P 2P lending. *Harvard Business Law Review* 6 (Online): 87–110. <http://www.hblr.org/2016/08/it-aint-broke-the-case-for-continued-sec-regulation-of-p2p-lending/>.
- Lo, Benjamin. 2016b. Online lenders shouldn't get mad over madden. *The Journal of Business, Entrepreneurship & the Law* 10. HeinOnline: 63.
- Lufax. 2015. *P2P lending market in China*. <http://blog.lendit.com/wp-content/uploads/2015/04/Lufax-white-paper-Chinese-P2P-Market.pdf>.
- Manbeck, Peter, and Marc Franson. 2016. *The regulation of marketplace lending: A summary of the principal issues (2016 Update)*.
- Manbeck, Peter, and Marc Franson. 2018. The regulation of marketplace lending: A summary of the principal issues. https://www.chapman.com/insights-publications-Regulation_Marketplace_Lending.html.
- Megaw, Nicholas. 2018. UK P2P Lender Looks into Selling the Loans on Its Platform. *Financial Times*, October 29, 2018. <https://www.ft.com/content/a19cd1d8-db65-11e8-9f04-38d397e6661c>.
- Milne, Alistair. 2018. Marketplace lending: Business models and regulation in australia and the UK.
- Milne, Alistair, and Paul Parboteeah. 2016. *The business models and economics of peer-to-peer lending*. 2016, no 17. Research Report. Centre for European Policy Studies. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2763682.
- Nahata, Pallavi. 2018. *Regulations shrink India's peer-to-peer lending industry*. August 18, 2018. <https://www.bloomberquint.com/business/regulations-shrink-indias-peer-to-peer-lending-industry#gs.awnIMRI>.
- Neuwirth, Suzie. 2017. Wellesley puts P 2P product on 'pause' and boosts transparency. *P2P Finance News*, May 2017. <http://www.p2pfinancenews.co.uk/2017/05/26/wellesley-p2p-pause-transparency-results/>.
- People's Bank Of China. 2015. *Guiding opinions on advancing the healthy development of internet finance*. Beijing. http://www.gov.cn/xinwen/2015-07/18/content_2899360.htm.
- Perkins, David W. 2018. Marketplace lending: Fintech in consumer and small-business.
- Pranjivan, Vinay. 2017. *Financial innovation and alternative finance: A comparative analysis of the objectives of regulation and its impact on lending based crowdfunding in France and in the UK*.
- P2P Research Group. 2016. *P2P with Chinese characteristics: Development, regulation and outlook*. <http://www.cf40.org.cn/uploads/research/20160115p2p.pdf>.
- Samitsu, Atsushi. 2017. *Structure of P2P lending and investor protection: Analyses based on an international comparison of legal arrangements*. Bank of Japan.

- Shen, Wei. 2015. Internet lending in China: Status quo, potential risks and Regulatory options. *Computer Law & Security Review* 31: 800.
- Shen, Wei. 2016. *Shadow banking in China: Risk, regulation and policy*. Edward Elgar.
- Wardrop, Robert, and Tania Ziegler. 2016. A case of regulatory evolution-a review of the UK financial conduct authority's approach to crowdfunding. *DICE Report* 14 (2). CES-IFO GmbH: 23.
- Weinland, D., and Ju S. F. 2018. China's P2P lenders say regulation will cause industry collapse. *Financial Times*. <https://www.ft.com/content/eac2c2de-d050-11e8-a9f2-7574db66bcd5>.
- Wildau, Gabriel, and Yizhen Jia. 2018. Collapse of peer-to-peer lenders sparks investors flight. *Financial Times*, June 22, 2018. <https://www.ft.com/content/75e75628-8b27-11e8-bf9e-8771d5404543>.
- Wilmarth Jr., Arthur E. 2004. OCC's preemption rules exceed the agency's authority and present a serious threat to the dual banking system and consumer protection. *Annual Review of Banking and Financial Law* 23. HeinOnline: 225.



Regulating Equity Crowdfunding: Lessons from International Practice

Emmanuel Schizas, Thomas Kenneth Bennett,
and Waleed Samarah

1 Part 1: Introduction

Equity crowdfunding is the activity of facilitating small-scale offerings of equity in unlisted firms by matching the fundraiser to members of the public via an online platform. The bulk of the sector is dedicated to funding seed- and early-stage businesses, but a significant share deals in the equity of special purpose vehicles (SPVs) invested in real estate and other illiquid assets. While this sector technically came into existence in 2011 with the founding of the UK platform Crowdcube, it draws on a rich history of analog-era predecessors, from business angel networks and syndicates to early stock exchanges and even today's junior stock exchange boards.

Estimates compiled from multiple CCAF reports (World Bank and CCAF Forthcoming) suggest that the global equity crowdfunding industry channelled more than \$1.38 bn into early stage firms in 2017 down from a peak of just under \$3 bn in 2015, and more than \$3 bn into crowdfunded property projects—up from \$676 m respectively in 2015. The USA, United Kingdom and China have historically led the sector's growth, although the outsize dominance that Chinese platforms have over global P2P lending is not observable in the case of equity crowdfunding. In fact, since 2016 Chinese

E. Schizas (✉) · T. K. Bennett · W. Samarah
Cambridge Judge Business School, University of Cambridge, Cambridge, UK
e-mail: e.schizas@jbs.cam.ac.uk

equity crowdfunding, in the strict sense of the word, has retreated significantly as a result of regulatory intervention (see Part 6). In relation to real estate crowdfunding, the very high ranking of South Korea (as the 3rd largest market globally, behind the USA and United Kingdom) is also worth noting.

The figures above make equity crowdfunding a very small sub-set of seed- and early-stage equity financing globally, which topped \$80 bn in 2017 (Rowley 2019). Yet in some of the more advanced markets equity crowdfunding has evolved into a significant part of the seed- and early-stage funding ecosystem—most notably in the case of the United Kingdom where the crowd accounts for 13% of all such funding and about a third of all business angels have co-invested with crowd investors (Zhang et al. 2018). Importantly, regulation is one of very few plausible institutional determinants of equity crowdfunding growth. Rau (2019), controlling for the influence of a wide range of explanatory variables, finds a small but statistically significant correlation between equity crowdfunding as a share of total crowdfunding volume in a given jurisdiction and the perceived quality of (generic) regulation in the same jurisdiction.

This chapter explores further the way in which equity crowdfunding is regulated in practice, encompassing both rule-makers' and supervisors' perspectives. It is substantially informed by two core data sources. The first is the inaugural Global Alternative Finance Regulation Survey, jointly conducted by the World Bank and the Cambridge Centre for Alternative Finance, which received 99 responses from regulators representing 111 jurisdictions. The headline findings of that survey are presented in (Rowan et al 2019). The second source is the 2016 Equity Crowdfunding Investor Survey carried out by the CCAF on behalf of the UK's Financial Conduct Authority (FCA & CCAF, unpublished). That survey was carried between 11 November 2016 and 8 January 2017 and achieved 1,379 responses from investors on UK Equity Crowdfunding platforms. To the authors' knowledge, this is the largest ever survey of such investors, and a summary of the survey methodology and findings is reported in Zhang et al. (2017).

2 Part 2: The Risk Profile of Equity Crowdfunding as a Function of the Sector's Value Proposition

2.1 Democratising Venture Capital and the Crowd's Place in the Pecking Order

Early stage investment involves high risks and (potentially) high returns: most startups fail and success stories returning high multiples of the initial investors' money are rare. The typical business angel investment thus tends to result in a net loss (Kisseleva 2019; Wiltbank 2009), and there is rarely a secondary market into which such assets can be sold if the investor needs liquidity urgently. Early stage investors have, thus, usually relied on building diversified portfolios, such that the minority of very successful investments can make up for losses elsewhere and still provide average annualised returns in the order of 20–40% (see e.g. Gregson et al. 2017).

The diversification strategy of early stage investors (taking many small losses for a few big wins) has traditionally been available only to high-net-worth individuals and professional or institutional investors. Depending on the method of estimation, between 10 and 50 separate investments might be needed for an early stage portfolio to become reasonably diversified (Gregson et al., *op cit.*). Because investors and fundraisers need to commit resources to any one transaction (e.g. for due diligence and legal advice), small individual investments aren't worthwhile for either party. In the United Kingdom, where the typical angel investment is in the order of £150,000 (Wiltbank, *op cit.*), properly diversified investors are likely to have committed seven-digit sums to the asset class.¹ Inevitably, this has traditionally limited the pool of funds from which even the best-placed early stage businesses could hope to be funded.

Crowdfunding makes this same investment strategy available to individual, non-professional or retail investors by creating an efficient way for them to take much smaller stakes in the equity of multiple companies. To do this, such investors rely on issuer disclosures, a degree of due diligence by the platform operators and further due diligence performed by fellow investors—including institutions. Accordingly investors prioritise perceived openness and transparency in their evaluation of platforms (Samarah and Alkhatib 2019).

¹ See also Oxera (2015). Oxera make the same argument assuming a much smaller minimal investment of £10,000.

Crowd investors do not usually invest on the same terms as founders, business angels or institutional investors. They rarely have an opportunity to negotiate company valuations, and often buy non-voting shares which give them little to no ongoing role in the governance of the business. All of these compromises to investor power, where they occur, are associated with higher fundraiser risk profiles (Cumming et al. 2019), as compared to investment in so-called A-shares. Where retail or individual investors form the greater part of the 'crowd', it is in any case common for companies to only offer between 10 and 15% of their equity to such investors (Cumming and Johan 2019). Furthermore, crowdfunding platforms themselves do not undertake to act in the interests of crowd investors following a funding round, even when they offer to hold shares under a nominee structure. Thus, despite exceptions, crowd investors are best thought of as passive investors (Hornuf et al. 2019).

Pecking-order theory predicts that equity crowdfunding should suffer from adverse selection, in the sense that firms might turn to the crowd only if they are unable to obtain funding otherwise, and if they believe that crowd investors will accept valuations that venture capitalists will not. The extreme information asymmetry involved in early stage investment could make this a rewarding strategy for lower quality firms. Walthoff-Borm et al. (2018a, 2018b) provide empirical evidence of adverse selection using matched studies of UK ECF and non-ECF funded firms. These suggest that the failure rate of would-be ECF fundraisers is higher than that of comparable firms, and that the former tend to be less profitable and more indebted. Blaseg et al. (2020) find tentative evidence of the same in Germany. Such findings are also in line with the experience of investors in unlisted equity outside the crowdfunding sector. In the USA, investment in unlisted, over-the-counter shares has likewise historically been associated with very poor investor outcomes, particularly where the issuer has operated an unregulated promotional campaign (White 2016).

There is only limited data to date on the actual performance of crowd-funded startups, due to the small number of exits and the private nature of many follow-on funding rounds. Formula-driven valuations, e.g. from Altfi (2016) for UK platforms, suggest double-digit annual returns across the entire population, after accounting for conditional fiscal incentives. Returns net of incentives both in Altfi (2016) and in more limited studies using formula-driven valuations, such as Signori and Vismara (2016) range between 8 and 9%. Such returns are, at first glance, inferior to those that a business angel might expect (Wiltbank 2009). However paper portfolio returns calculated on very recent investments without an exit can be misleading: Wiltbank (2009) suggests that it takes about three years for investors to exit a failed

investment but about seven years for them to exit a successful one. This should imply that, where the sector is still young, returns across the universe of funded firms might appear artificially low.

2.2 The Value and Limits of Outsourced Due Diligence

Based on the sector's value proposition, it seems fair to say that the social value of equity crowdfunding depends on the ability of crowd investors to create properly diversified portfolios, the ability of issuer disclosures to facilitate the accurate pricing of securities, and the cost-effectiveness of outsourcing due diligence to unregulated fellow investors and the platforms themselves as a passive investor. Regulation can in principle influence all of the above.

Price discovery, as understood in formal capital markets, is unlikely to take place on most equity crowdfunding platforms currently in operation, but merely approximated. Crowd investors are known to use simple proxies of issuer quality that are typically easy to obtain in a hosted pitch scenario—Mamonov and Malaga (2018) provide a helpful review, citing, e.g. the provision of financial projections, the degree of equity retention, the presence of patents, the depth of founders' social networks, and the presence of well-known investors among the issuer's shareholders. Åstebro et al. (2018) add to this list a theoretical framework of herding in equity crowdfunding markets that anticipates observed behaviour—and suggests that investors use the size of recent pledges and the degree to which a pitch has been funded as proxies of issuer quality. Such herding might be rational if investors are sensitive to two types of signals: pitch momentum might indicate that knowledgeable and trusted individuals in 'the crowd' have vetted the pitch and are convinced of its merits; while a large number of small contributions might be a good proxy for consumer demand for the fundraiser's products, or the firm's ability to generate word of mouth. Crowd investors in other settings, such as rewards-based crowdfunding, are known to reward robust project risk disclosures and to use them in their due diligence (Madsen and McMullin 2018); however these are typically absent from equity crowdfunding, and must be reproduced at great cost, if at all.

In summary, some investors outsource part of their due diligence to a third party in one of two ways: a 'validation route' relying on the reputational capital of firms and individuals, and a 'herding route' relying on pitch momentum. Empirical evidence suggests that the degree of outsourcing can be very significant. Findings from the FCA & CCAF (unpublished, op cit) suggest that 56% of UK equity crowdfunding investors considered the

amount already pledged to a fundraiser to be ‘important’ or ‘very important;’ whereas 57% valued the participation of an experienced investor in the same way. The median investor reported spending about 40 minutes per week on due diligence on equity investments,² contrasting with Wiltbank’s (2009) threshold of 20 hours’ due diligence per deal. 76% of respondents to the FCA and CCAF survey agreed or strongly agreed that they relied on due diligence carried out by the platform operator when selecting investments; 25% strongly agreed.

Each of these findings is potentially problematic from a regulator’s perspective. Reliance on proxies leaves investors more vulnerable to exaggerated claims by malicious issuers. A reliance on external validation can expose the platform and its larger investors to conflicts of interest, and investors and platform operators may have very different perceptions of what constitutes ‘due diligence’ unless this is set out in regulations. Finally, the herding route offers no guarantee that due diligence is being carried out at all, and can be gamed, e.g. by directing pre-pledged investment through the platform.

2.3 Can Effective Diversification Be Facilitated on Crowdfunding Platforms?

It is not clear that on-platform diversification is an effective mitigant of risk in practice. As of 2016, the median UK equity crowd investor had only invested in about five businesses (CCAF and FCA, unpublished). Only the top 7% of investors had stakes in more than 10 businesses. For the purposes of Gregson et al. (op cit), crowd investors would thus be very poorly diversified, and heavy off-platform diversification would be needed to make up for the unbalanced nature of their portfolios. Yet the CCAF and FCA (unpublished) survey findings suggest that about 20% of UK equity crowd investors (typically those with smaller exposures) did not, as of late 2016, consider off-platform diversification to be an important element of their strategy.

More importantly, equity crowdfunding portfolios may be structurally biased regardless of size. It is well-documented that early stage investors are more comfortable investing in industries they know and products they can understand. Crowd investors, however, need not have any specialist industry experience. As a result, some crowdfunding platforms may over time have evolved a specialism in fundraising for businesses that serve consumer needs familiar to the niche demographic of crowdfunding investors (see e.g.

² This statistic was derived by obtaining midpoint values for 7 ranges and should therefore be treated as indicative only.

Lukkarinen et al. 2016). Typical examples are consumer electronics and smart accessories, fintechs, premium food and drink, or sharing economy applications. If this effect is significant and persistent, then a portfolio made up of very many crowdfunded businesses could be poorly diversified due to exposure to common underlying factors—e.g. disposable incomes of the global, professional middle or upper-middle class.

3 Part 3: Operating Models of Alternative Fundraising

Regulatory frameworks for equity crowdfunding have often been developed against the backdrop of a rapidly evolving sector, where it is hard for regulators to know whether observed risks are inherent to the untested business models in the market or idiosyncratic to the (often young) firms they regulate. In this section, we propose a simple typology of equity crowdfunding business models that might allow for a more structured and nuanced discussion of risks and market outcomes in this sector. We distinguish between a basic ‘Hosted Pitch’ model; a hybrid ‘Syndicate’ model; versions of the two that include crucial complex elements such as secondary markets or funds; and the materially distinct ‘Real Estate’ model.

3.1 The Hosted Pitch

Under this model, a financial intermediary facilitates direct transactions in securities via an online platform and provides tools and advice for maximising the reach of the issuer’s crowdfunding campaign. The platform operator typically carries out a regulated activity—for example, approving financial promotions, arranging transactions in securities or similar.

In the hosted model, the platform operator takes responsibility for checking factual claims, but not for devising or challenging proposed valuations and target fundraising amounts. In fact, most hosted pitches are tightly controlled in order to limit the host’s liability. For example, platforms might not include any financial information or forecasts in the hosted pitch, but facilitate direct messaging between fundraisers and investors who want this information.

3.2 The Syndicate Model

Under the Syndicate or Hybrid Crowd model (see Chen et al. 2016 for an introduction) the issuer carries out a hosted pitch approved by the platform operator, but deals with the crowd through a lead investor: a member of the platform who is an experienced business angel and has agreed to take a stake in the business after carrying out their own due diligence. Unlike with a hosted pitch, the fundraiser's investment terms are likely to be negotiable, with the lead investor questioning, e.g. the valuation of a company. The Syndicate model is, in this regard, similar to the way many business angel networks have traditionally worked.

3.3 Complex Structures: Secondary Markets and Funds

More complex fundraising platforms with a large back-book of funded firms have been known to operate **secondary markets** for those firms' securities in return for fees on any transaction. In practice, participants in such markets are unlikely to be able to trade as they wish, with platform operators setting restrictive rules in order to manage their own liability and compliance burden—this may include rules limiting price discovery or the provision of liquidity. Fully functional secondary markets for illiquid securities are likely to be much more heavily regulated than mere promotions, with certain obligations extending to issuers as well as platform operators. For example, in the EU, such platforms are likely to be subject to regulation relating to Multilateral Trading Facilities (MTFs) (see ESMA 2014) while in North America they might be regulated as Alternative Trading Systems (ATS), which also requires licencing as a broker-dealer.³ Platform operators wishing to provide investors with exit opportunities may alternatively, and subject to additional licences, operate **funds** that invest in their own crowdfunded securities.

3.4 Real Estate Crowdfunding

Under the real estate crowdfunding model, the platform operator offers investments in illiquid assets—typically but not exclusively real estate—which it originates by itself, e.g. using an affiliate or referral network. For each investment opportunity, investors' funds are pooled in a special purpose

³ The first examples of this are US platforms SeedInvest and StartEngine.

vehicle (SPV), notionally independent of the platform operator, which creates or buys the target assets. The investors, in turn, are given equity stakes in the SPV. Real estate crowdfunding is thus a significant departure from other crowdfunding segments in that it is not a two-sided market.

In many jurisdictions, the regulatory status of the property crowdfunding model is complicated by the presence of an SPV, as firms have been known to abuse such structures. For example, the UK's FCA treats such offerings as 'Non-Mainstream Pooled Investments' (NMPIs) and does not generally allow firms to market them to ordinary retail investors without a waiver from the regulator.⁴ Funds investing in similarly non-standard and illiquid assets are typically also restricted from marketing to the mass market in many jurisdictions—this is true, for example, of Unregulated Collective Investment Schemes (UCIS) under UK law, for many Alternative Investment Funds (AIFs) across the EU, and for vehicles whose securities are offered through Regulation D placements in the USA.

4 Part 4: An Equity Crowdfunding Risk Register

In their survey of regulators Rowan et al. (2019) found that governments and regulators see significant promise in alternative financing methods such as equity crowdfunding, and particularly strong benefits to SME financing and competition among funding providers. However these must be balanced against the risks deducible from the literature discussed above, and summarised in Table 1.

The findings of Rowan et al. (2019) reveal major differences in perception between regulators with supervisory powers over equity crowdfunding platforms and those with no such remit. The former are more likely to cite mundane capital and liquidity risks as their top priorities, as compared to non-supervisors who were more likely to cite money laundering, fraud and cyber-security risks. This might imply that, from the perspective of active supervisors, equity crowdfunding is a fairly traditional financial services activity, and its tech-enabled aspects are only incidental to its risk profile.

As with any high-risk investment, capital losses in this sector are not in themselves evidence of harm. In fact, the CCAF and FCA (unpublished, op cit) found that investors saw equity crowdfunding as by far the riskiest of the alternative finance asset classes. 71% of this sample saw their

⁴ See e.g. FCA (2018).

Table 1 Risks to equity crowdfunding Investors

Expected risks	% of actively supervising regulators citing in Rowan et al. (2019)	Conditions for risks to be of particular interest to regulators (authors' assessment)
High rates of capital loss and/or dilution	82%	<ul style="list-style-type: none"> • Adverse selection; failure rate higher than other early stage investment; • Investors consistently underestimating the probability of capital loss • Investors highly exposed and under-diversified • Consistent mispricing of equities/company assets or downward revaluations as part of follow-on funding rounds
Fraud, misuse of funds	66% (fraud) 37% (misuse for criminal purposes)	<ul style="list-style-type: none"> • Evidence of 'rolling bad apples', i.e. repeat failed fundraisers • Highly unusual or opaque SPV structures
Severe liquidity risks	61%	<ul style="list-style-type: none"> • Investors consistently overestimating the liquidity of assets; • Investors with a short-term investment horizon exposed to the sector
Conflicts of interest (principal-agent and principal-principal)	Not prompted	<ul style="list-style-type: none"> • Retail/ordinary investors investing simultaneously but on different terms than VCs/other backers (B- and non-voting shares) • Investors assuming that firms will act in their interests post-issuance • Investors assuming that intermediaries are liable when they are not (e.g. in relation to due diligence conducted) • Platform financial incentives aligned to origination, not performance

crowd equity exposures ‘mostly’ or ‘completely’ as speculative investment, and on average they saw the asset class as having a loss probability similar to forex trading (though less than crypto-currencies).⁵ Arguably, where investors have ex ante expectations of this kind, and are additionally fairly well-off and sophisticated, fewer adverse outcomes will represent genuine harm to them. It is also possible to argue that regulators faced with such an investor population should develop a higher risk appetite in relation to equity crowdfunding, allowing them to redirect supervisory resources to sectors with more risk-averse and potentially vulnerable investors.

This however has not been the case historically. Rowan et al. find that enforcement against equity crowdfunding platforms is common around the world—across all of the markets for which data were provided, the total number of historical cases was between 10% and 20% of the population of currently active firms.⁶ This implies that risks related to fraud or platform misuse or breaches of securities law may in fact be higher organisational priorities than the findings in Fig. 1 might suggest.

Particular challenges are posed by the potential for crowdfunding platforms to facilitate cross border financing. The Financial Stability Board (FSB) has emphasised the need for international cooperation between regulators, focusing particularly on the need to manage third-party operational risks, to mitigate cyber risks by sharing information, and to monitor macro-financial risks by improving the capability of the regulators to access sources of information and data (Financial Stability Board 2017).

5 Part 5: The Regulatory Response to Equity Crowdfunding

Based on responses from 111 jurisdictions around the world, Rowan et al. (2019) provide a highly standardised description of how equity crowdfunding is regulated around the world. At the highest level, regulatory frameworks for crowdfunding platforms, including equity crowdfunding, can be classified as either reactive, if regulators rely on pre-existing regulations (typically

⁵ It is worth noting that investors were not asked to rank different crowdfunding asset classes in terms of risk. Rather, over investors in each alternative finance model (9,000 across all models) were asked to compare the riskiness of the model in question relative to other potential investments, such as AAA-government bonds, cash in the bank, buy-to-let property, blue chip equities, corporate bonds, widely traded currency pairs, and cryptocurrencies.

⁶ A precise estimate (18%) is available on the basis of the CCAF and World Bank data. However, because firm numbers and enforcement cases were both reported as ranges and subsequently converted to midpoint estimates for further analysis, it is preferable not to rely on a single derived ratio.



Fig. 1 Share of regulators in Rowan et al. (2019) citing different risks among their top priorities when supervising Equity Crowdfunding—by formal remit over the sector

as securities law), or proactive/bespoke, if tailor-made regulations and regulatory perimeters are developed for the sector (Aggarwal and Stein 2016). According to Rowan et al., nearly four out of ten jurisdictions globally regulated the sector as of mid-2019, and that this percentage might rise to over two thirds by 2021; 22% of all jurisdictions employed some form of proactive or bespoke regulation, and an additional 8% had created sector-specific adjustments or exemptions to pre-existing regulation to accommodate equity crowdfunding. More than half of the purpose-built regulatory frameworks developed for this sector were created between 2015 and 2017.

Equity crowdfunding differs from other alternative finance activities in that it has in most cases come into existence as an already-regulated sector, and, unlike P2P lending or initial coin offerings (ICOs), whether or not a given regulator had remit over it has often been clear and uncontroversial. This is because equity crowdfunding typically involves the marketing of securities, which tends to be a regulated activity under long-established securities law in most jurisdictions. The authorities charged with supervising this sector were in almost all cases (95%) securities regulators, but there was also evidence that fund management, wholesale banking and marketing of financial products were areas of regulatory expertise that applied themselves readily to regulating equity crowdfunding.⁷

⁷ This finding was obtained by comparing the sector remits of respondents that had responsibility for equity crowdfunding (40% of the sample) with the remits of respondents that did not.

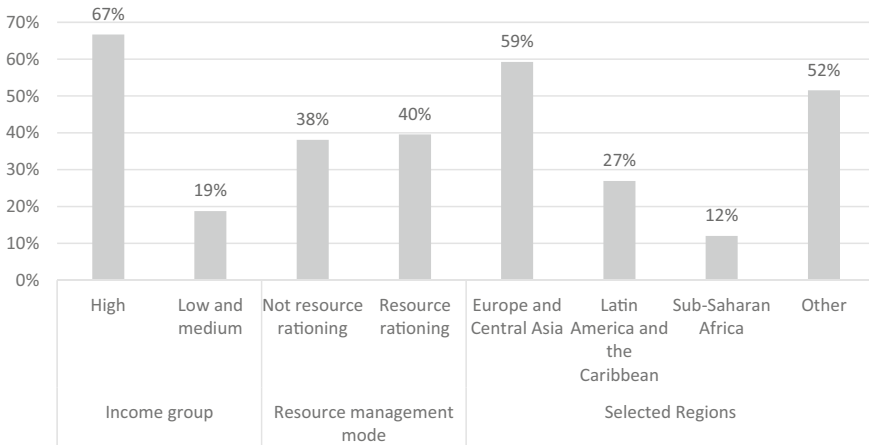


Fig. 2 Incidence of alternative finance regulation in Rowan et al. (2019)

As Fig. 2 shows, equity crowdfunding was, as of mid-2019, still much more likely to be regulated in high-income countries,⁸ and was particularly unlikely to be regulated in Sub-Saharan Africa. In developing countries, the low prevalence of regulation can be explained by limited industry growth, a low level of awareness among the main stakeholders and limited financial literacy among both potential funders and fundraisers. However, the sequence implied above might just as easily be reversed; Samarah and Alkhatib (2019) suggest that the development of a proportionate regulatory framework might be one highly effective means of raising awareness of crowdfunding, and Baeck et al. (2014) found that businesses would be much more willing to use equity crowdfunding if it were regulated.

Of the three fintech verticals studied (P2P or marketplace lending, equity crowdfunding and initial coin offerings), equity crowdfunding emerged as the most proactive in welcoming and shaping regulation. In fact, 67% of regulators that had recently started or completed a review of regulations in this area cited proactive calls for regulatory change from industry as one of the triggers for this. It is likely that the close proximity to the perimeter for securities regulation forces platform operators in this sector to seek legal and regulatory certainty more urgently than other firms.

⁸ World Bank Country and Lending Groups <https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups>. [Accessed: 31 May 2019] Note that 'resource rationing' jurisdictions are jurisdictions in which the responding regulator claimed that resource limitations made it harder, all other things being equal, to supervise alternative financing sectors such as equity crowdfunding than traditional activities of this kind. Rowan et al. (2019) use this as a proxy for the regulators' strictness in rationing their resources, and find that it does not correlate significantly with country income level.

As with all sectors studied by Rowan et al. (2019), the purpose of equity crowdfunding regulation appears to be to establish a reasonable boundary between investments that are appropriate for individual, non-professional investors and those that are not, and to create an acceptably safe market for those investors that they see as legitimate users of the platforms. As a result, most jurisdictions in which the sector was regulated allowed for promotions to individual investors—however only about half allowed firms to hold client money in the process of facilitating transactions. Few jurisdictions anticipated permissions to operate a secondary market, and in practice those that did had very little such activity in place to speak of.

In their review of the regulatory frameworks for crowdfunding in six countries, EBRD and Clifford Chance (2018) suggest that crowdfunding regulation ought to provide for (i) continuity and ongoing administration of platforms (ii) capital and liquidity requirements; (iii) Know Your Customer (KYC) rules and anti-money laundering (AML) checks; (iv) a cap on the size of offerings; (v) caps on the amount a single investor may commit to the platforms; (vi) consumer protection measures, including investor disclosures and risk warnings; and (vii) due diligence and pre-funding checks.

While it is possible to argue whether or not EBRD and Clifford Chance (2018) set out a definitive standard of good practice, it is clear that few of the elements they argue for are consistently applied across jurisdictions. As of mid-2019, all regulatory frameworks for equity crowdfunding relied heavily on firm disclosures rules (see Table 2). All jurisdictions with regulated equity crowdfunding markets had high-level rules in place requiring honest disclosures and almost all prescribed some standard types of issuer information that ought to be provided in all cases. However, beyond these high-level similarities there is no consensus on the right approach to regulating this sector.

Instead regulatory frameworks in this area are best distinguished by contrasting their responses to the following questions:

- Is an equity crowdfunding round a public offering of securities?
- Is the platform operator providing an investment or brokerage service?
- How desirable is it that non-professional investors be able to access the underlying asset class?

This approach requires a recasting of the Rowan et al. (2019) dataset to focus on four alternative models of regulation for this sector: the **gate-keeper**, **pure intermediary**, **bespoke hosted** and **public offering** models.

This section, including Tables 2 through 4, draws on this new classification to more clearly illustrate the nature of each regulatory approach (Table 3).

5.1 The Gatekeeper Model

Under this model, platforms are expected to facilitate a market for professional, high-net worth or highly knowledgeable investors only. It is a fairly common model—21% of jurisdictions that regulate equity crowdfunding did not, as of mid-2019, allow crowdfunding platform operators to promote securities to individual, non-professional investors at all. Such jurisdictions include the likes of Estonia, Argentina, Greece, Lebanon or the Qatar Financial Centre, and the model is particularly common in non-OECD jurisdictions. Notably, although China's regulatory framework for equity crowdfunding formally anticipates marketing to the mass market ('public equity crowdfunding'), the lack of approved regulations for this activity has made public equity crowdfunding effectively illegal since 2016, and has forced platforms to migrate to a de facto gatekeeper model. In the USA, the gatekeeper model is represented by Title II equity crowdfunding, which is restricted to accredited investors, whereas in the United Kingdom certain property crowdfunding offerings employing an SPV (non-mainstream pooled investments (NMPs)), are not generally allowed to be marketed to non-professional, or retail, investors.

It is likely that in excluding individual investors the Gatekeeper model aims to justify a more lightly regulated market—but, as Table 4 shows, the aim is not an unregulated market either. It is likely that such regulatory frameworks are motivated by a perception of the asset class as a highly risky one, as many still aim to offer significant, costly and even paternalistic protections to professional investors. Gatekeeper sectors are thus possibly best seen as experimental complements to public equity markets, with regulators looking to build trust rather than drive transaction volumes.

5.2 Pure Intermediary Model

Under the pure intermediary model, the platform operator must be an authorised firm, but it is treated as having very little input beyond putting eligible investors and issuers in touch with one another and executing transfers of funds between them. As a simple proxy for this model, just under one in five jurisdictions (19%) with a regulatory framework in place do not require platform operators to verify claims made by issuers hosted on their platforms.

Countries such as Belgium, New Zealand or Brazil have taken this light-touch approach, and it appears to be most common among smaller OECD countries. While the pure intermediary model might impose fewer ongoing requirements, a stricter authorisation process or management responsibility regime might in some cases be employed to compensate for this; investor exposure caps or, as in Belgium, exposure caps for all but the smallest funding rounds, might assist further in this direction. Elsewhere, e.g. in the case of New Zealand or Brazil, the potential risks of lighter regulation might be mitigated by a market or regulatory bias towards a Hybrid Crowd or Syndicate model (see e.g. Schwartz 2018 in the case of New Zealand).

As Table 2 shows, the **pure intermediary model** is light on onboarding requirements, and to some extent on disclosures, but it is not lightly regulated across the board; systems and controls and anti-money laundering requirements are no less common here than with other models, and firms are usually subject to client money regimes. Pure intermediary regulation thus appears to be an attempt to bring into the regulatory perimeter what might otherwise be unregulated equity investment bulletin boards and digitally enabled business angel networks.

5.3 Bespoke Hosted Model

Under the most common regulatory approach, favoured by about 40% of all jurisdictions, crowdfunding platforms are treated as carrying out a distinct type of intermediation and are expected to take a certain measure of responsibility not just for promoting securities to the general public but throughout the lifecycle of a transaction, including vetting promotions and carrying out due diligence, or taking responsibility for client money. Jurisdictions such as Abu Dhabi, Australia, France or Mexico have taken this approach, as did the United Kingdom with its regulatory framework of 2014. Importantly, regulators benchmark against each other and against international good practices before developing new rules, and bespoke hosted frameworks are twice as common among the top 10 most-benchmarked against jurisdictions⁹ as they are among other jurisdictions. Hence it is likely that this model will continue to gain in popularity.

⁹ Rowan et al. find (op cit) that these include the UK, USA, Singapore, Spain, France, Malaysia, Australia, UAE, New Zealand and Mexico.

Table 2 Regulatory requirements in regulated equity crowdfunding markets, by model of regulation (Rowan et al 2019)

Detail of requirement	Type of regulatory requirement	Total N = 43 (%)	Gatekeeper N = 9 (%)	Public offering model N = 9 (%)	Bespoke hosted model N = 17 (%)	Pure intermediary N = 8 (%)
General requirement that communications with customers are accurate and complete	Promotions and disclosures requirements	100	100	100	100	100
Requirement to provide standardised information to investors (e.g. risk warnings, costs, incentive structures, etc.)		91	100	100	93	75
Requirement to promote to investors differently based on wealth and/or experience		47	60	71	29	50
Requirements (such as restrictions or conditions which apply to advertising using specific types of media (e.g. social media))		29	20	14	43	25
Anti-Money Laundering requirements (e.g. KYC) specific to this type of firm	Onboarding requirements	88	100	86	86	88
Requirement for the firm (or a third party) to verify the claims made by the fundraiser		74	80	100	100	0

(continued)

Table 2 (continued)

Detail of requirement	Type of regulatory requirement	Total N = 43 (%)	Gatekeeper N = 9 (%)	Public offering model N = 9 (%)	Bespoke hosted model N = 17 (%)	Pure intermediary N = 8 (%)
Requirement to ensure eligibility criteria are met before fundraising (e.g. minimum company age, profitability)		82	80	100	86	63
Requirement to provide fundraisers/investors with cancellation rights		59	60	71	64	38
Requirement to cap the amount invested by an individual at a defined percentage of their wealth or income		53	80	71	50	25
Mandatory governance requirements (e.g. independent risk management, internal audit)	Systems and Controls, Risk Management and Prudential Requirements	85	100	100	79	75
Mandatory complaints handling requirements		76	60	71	93	63

Detail of requirement	Type of regulatory requirement	Total N = 43 (%)	Gatekeeper N = 9 (%)	Public offering model N = 9 (%)	Bespoke hosted model N = 17 (%)	Pure intermediary N = 8 (%)
Requirement to have a 'wind-down' plan to minimise impact of platform failure (e.g. resolution packs or living wills)		53	20	71	57	50
Requirement to separate client assets from firm assets		82	100	86	86	63
Requirement to share relevant user data (e.g. with a public or private credit registry)		21	0	43	21	13
Requirement to hold capital proportionate to total amount invested via the platform		21	20	43	0	38

Bespoke regulatory frameworks for equity crowdfunding are, as the name suggests, developed purely for this activity, working with and around pre-existing securities law. To this end, some frameworks include broad exemptions from the status of a public offering (as is the case, e.g. in Abu Dhabi) or highly tailored offering types with lighter transparency requirements (as is the case, e.g. in Australia). As with ‘pure intermediary’ regulatory frameworks, bespoke frameworks intentionally allow the sector to reach a mass market. They may therefore employ investor exposure caps (either in absolute terms or as percentages of individual income or wealth),¹⁰ in order to reduce the overall risk to individual investors. This is the approach taken, for example, in the United Kingdom or Australia. Because firms must obtain licences specific to this sector, regulators have additional options to tailor the licencing process to the types of risks most prevalent in the sector.

Bespoke frameworks are perhaps best understood as representing the trade-offs made by regulators who are driven not to simply tolerate and sanitise but actually increase retail participation in this market. Rowan et al. (2019) remark on the relative absence of a strict individual vs professional investor distinction among many bespoke regulatory frameworks, intended to ensure that retail investors can participate in the market. Nevertheless, some bespoke frameworks incorporate softer restrictions that still aim to filter out less knowledgeable or experienced investors; the UK regime, for example, includes a requirement that firms assess how well non-advised individual investors understand the asset class, in what is known as the *appropriateness* test.¹¹

In line with the objective of making equity crowdfunding a retail market proposition, as Table 2 shows, there is a greater emphasis on creating explicit due diligence requirements, online and social media marketing standards, and complaints handling rules for the sector.

¹⁰ A cap of this nature might be ‘hard’, in that a firm must obtain evidence of the client’s financial position before onboarding them and exposing them to promotions, but in most cases it is ‘soft’, in that clients are relied upon to attest that they will not exceed the exposure level prescribed in regulation.

¹¹ See FCA COBS. Note that what is being assessed for ‘appropriateness’ is the act of facilitating a non-advised transaction in this asset class for the particular investor; the appropriateness test does not consider whether the investments they are entering into are appropriate to their needs.

5.4 Public Offering Model

Under the public offering model, employed by just over one in five jurisdictions (21%),¹² an equity crowdfunding round is treated as a new type of public offering of securities. While there may be exemptions or tailored requirements in relation to public disclosures and other obligations of the issuer, pre-existing securities law remains the point of departure, platform operators are typically regulated as brokers, and issuers are expected to file substantial regulatory returns even if they do not prepare full prospectuses. Such tailored disclosures and reporting may only be an option for firms raising limited amounts of money, with the relevant cap provided in regulation.

As Table 2 shows, the public offering model stands out primarily for its reporting requirements (which extend to issuers as well as firms) and its capital requirements, and is the most highly regulated of the four models. Such regulatory frameworks may be driven to a great extent by the need to avoid regulatory arbitrage with formal listings and equity brokerage, while allowing the crowdfunding sector to survive and grow.

Perhaps the most explicit example of this model is the first of two equity crowdfunding routes available in Taiwan, whereby crowdfunding listings are hosted by the Taipei Stock Exchange's Go Incubation Board for Startup and Acceleration Firms (GISA). However South Africa, Austria, Norway and Kenya also follow variations of this approach, and the USA's approach to Title III equity crowdfunding is arguably similar. A close, if more temporary, alternative to the Gatekeeper approach might be to restrict the activities in question alongside the type of investor that can be engaged, and essentially allow the sector to operate in a 'test and learn' or Sandbox environment. For example, Saudi Arabia's Financial Technology Laboratory/Sandbox initiative, the first two trial financial technology licences were given to equity crowdfunding platforms Manafa Capital and Scopeer (Hurst, 2018).

¹² The classification of the UK is particularly difficult. CCAF & World Bank (2019) do not second-guess the self-classification of regulators into bespoke and pre-existing regimes. However, contrary to the official response by the FCA, the UK framework is best understood as a bespoke one. If this interpretation is taken, then percentages for the bespoke hosted model and the public offering model ought to be 42–19% respectively.

Table 3 Mitigation strategies for key risks in equity crowdfunding, by regulatory model

Expected risks	Conditions for risks to be of particular interest to regulators (authors' assessment)	Gatekeeper model	Public offering	Bespoke hosted model	Pure intermediary
High rates of capital loss and/or dilution	<ul style="list-style-type: none"> Adverse selection; failure rate higher than other early stage investment; Investors consistently underestimating the probability of capital loss Investors highly exposed and under-diversified Frequent downward revaluations as part of follow-on funding rounds Consistently optimistic valuations, esp. by non-independent parties 	<ul style="list-style-type: none"> Investor accreditation requirement High-level disclosure requirements Detailed standardised disclosures Due diligence requirements 	<ul style="list-style-type: none"> Detailed standardised disclosures Due diligence requirements Investor exposure caps 	<ul style="list-style-type: none"> Risk warnings Appropriateness testing Due diligence requirements Investor exposure caps 	<ul style="list-style-type: none"> Risk warnings Mandatory eligibility criteria
Fraud, misuse of funds and other operational risks	<ul style="list-style-type: none"> Evidence of 'rolling bad apples', i.e. repeat failed fundraisers Highly unusual or opaque SPV structures 	<ul style="list-style-type: none"> High-level disclosure requirements Governance requirements 	<ul style="list-style-type: none"> Detailed standardised disclosures Governance requirements Segregation, management of client money Wind-down planning requirements 	<ul style="list-style-type: none"> Segregation, management of client money Complaints handling rules Wind-down planning requirements 	<ul style="list-style-type: none"> Segregation, management of client money

Expected risks	Conditions for risks to be of particular interest to regulators (authors' assessment)	Gatekeeper model	Public offering	Bespoke hosted model	Pure intermediary
Severe liquidity risks	<ul style="list-style-type: none"> Investors consistently overestimating the liquidity of assets; Investors with a short-term investment horizon exposed to the sector 	<ul style="list-style-type: none"> Restricted investor participation 	<ul style="list-style-type: none"> Risk warnings Investor exposure caps 	<ul style="list-style-type: none"> Risk warnings, Appropriateness testing 	<ul style="list-style-type: none"> Risk warnings Risk warnings
Conflicts of interest (principal-agent and principal-principal)	<ul style="list-style-type: none"> Retail/ordinary investors investing simultaneously but on different terms than VCs/other backers (B- and non-voting shares) Investors assuming that firms will act in their interests post-issuance Investors assuming that intermediaries are liable when they are not (e.g. in relation to due diligence conducted) Platform financial incentives aligned to origination, not performance 	<ul style="list-style-type: none"> No provision unless fraudulent 	<ul style="list-style-type: none"> Corporate disclosure/prospectus requirements Corporate governance/free float requirements for issuers 	<ul style="list-style-type: none"> Investor exposure caps High-level systems and controls rules Risk management and governance requirements Complaints handling rules 	<ul style="list-style-type: none"> High-level systems and controls rules Eligibility criteria Complaints handling rules

Table 4 Summary statistics on regulatory requirements by regulatory model

	Gatekeeper	Pure intermediary	Bespoke hosted model	Public offering model	Total n =
	n = 9	n = 9	n = 17	n = 8	43
No. of permitted activities (out of 7)	1.3	3.8	3.1	3.3	2.9
No. of firm obligations (out of 17)	5.7	8.0	8.5	9.3	8.0

6 Part 6: The False Start of China's Public Crowdfunding Market

China's Equity Crowdfunding sector provides a dramatic example of tensions and shifts between regulatory models and its implications. The sector has been on a road to reform since the first regulations (Private Equity Crowdfunding Administrative Measures) were published, in pilot draft form, by the Securities Association of China (SAC) in 2014. That original regime set out licensing criteria and marketing/promotions restrictions for platforms such that offerings could not be said to be open to the public—namely a high threshold for qualified investor qualifications and a cap on the number of investors per funding round (200). This allowed the small, informationally opaque issuers on crowdfunding platforms to raise funds from wealthy investors while complying with Chinese securities law. While consistent with a gatekeeper model, this approach was not entirely deliberate but dictated by the inability to amend the relevant legislation in the short term (Xu and Ge 2017).

More than one hundred platforms entered the market following the publication of the draft Administrative Measures. Such growth was facilitated by the fact that the draft Administrative Measures allowed internet financing platforms to merely register with the China Securities Regulatory Commission (CSRC) rather than obtain a full financial services licence, thus providing a cost-effective means of entering the sector (BOP Consulting 2017).

By the end of 2015, a new consensus appeared to be emerging among China's regulators, with both official definitions and regulatory requirements moving away from references to private equity and stressing instead the 'open, public, small-sum' character of equity crowdfunding (Xu and Ge, op cit). In August 2015, the CSRC's *Announcement to Inspections of Internet Equity*

Financing Activities even formally allowed for the existence of a public equity crowdfunding sector—potentially in line with a public offering model or even a bespoke model of regulation. However the rest of the regulatory apparatus never moved in step with this new language. Both the Announcement and the Guidance to Promote the Healthy Development of Internet Finance (July 2015) required public equity crowdfunding platforms to be licensed as such by the CSRC, which never produced a licensing route for such platforms—creating a de facto ban on public equity crowdfunding (BOP Consulting, op cit).

This equilibrium, whereby platforms were attracted to the sector by relatively low regulatory barriers to entry and relied on gatekeeper practices to ensure compliance, was destabilised in 2016, with the combined publication of the CSRC's 2016 Implementation Plan of Special Rectification on the Risks of Equity Crowdfunding and the NIFA's Internet Finance Disclosure Regulations of 2016. The Implementation Plan focused narrowly on ending the misuse of equity crowdfunding platforms for the purpose of prohibited activities. These included raising private equity investment funds; conducting unauthorised public offerings of false advertising in the form of exaggerated fundraiser pitches; fraudulent issuance of shares and other financial instruments; misappropriation or criminal conversion of investor funds (Lin 2017). The Disclosure Regulations added mandatory transparency requirements in relation to platform activities.

This combination of thematic supervision and policy interventions led to an exodus of firms from the sector and there has been scant evidence of recovery since then—funding volumes fell by 85% between 2015 and 2017 alone (Ziegler et al. 2019). Those platforms that did not exit the sector altered their behaviour. Some now only promote to accredited investors in a de facto gatekeeper model, requiring the latter to register as members; some set limits on investment amounts and the number of investors per pitch. Platforms have abandoned the term 'equity crowdfunding platform' in favour of 'private equity financing platform'. As a result, there has been little sign of activity in the public equity crowdfunding sector in China. Instead, platforms are conducting equity financing activities that are in fact traditional private equity financing.

7 Part 7: Reflections on Equity Crowdfunding Regulation—Is There Another Way Forward?

A viable, value-adding equity crowdfunding sector is one that allows cost-effective diversification and provides an efficient, controlled way of outsourcing due diligence to more experienced investors. Facilitating such a market is a challenge for regulators, made more so by the peculiarities of investing in early stage ventures and the limited amount of resource regulators can dedicate to a relatively small sector. The variety of approaches taken to date strongly supports this.

Although equity crowdfunding investors tend to both expect and have a tolerance for extreme risk (at least in the United Kingdom), jurisdictions favouring a light-touch approach to equity crowdfunding are in the clear minority, and enforcement activity has historically been strong. Instead, bespoke frameworks appear to be emerging as the most common way of balancing retail investor access with the protections necessary for retail participation. This may be partly the result of regulatory benchmarking, but it also speaks to the versatility of such regulation.

Going forward, it is possible to pinpoint three important pivots that financial regulation might need to perform in order to keep up with likely trends in a growing sector:

First, secondary markets and fund-like structures could become a greater part of the equity crowdfunding ecosystem than they currently are. If this should prove to be the case, regulators would have to consider how best to regulate platforms who operate a true market in securities and those who act as principals rather than purely as agents. Some jurisdictions, including the USA and United Kingdom, are able to fall back on existing regulation in some of these areas (e.g. MTF rules in the case of the UK), but most emerging markets will need to develop these from scratch.

Second, it is notable that regulators have largely avoided creating bespoke regimes for real estate crowdfunding, despite the clear differences between this and more straightforward early stage equity models. It is possible that a combination of regulatory arbitrage and high-profile failures will prompt at least some regulators to rethink.

Third, bespoke regulatory frameworks are best seen as work-in-progress. As existing frameworks mature, post-implementation reviews will likely reveal flaws to be addressed. If these result primarily from regulatory arbitrage or extreme information asymmetries then it is likely that the public offering model of regulation will become more attractive to regulators.

As post-implementation review findings emerge, it is reasonable to ask whether regulation that focuses largely on formal disclosures, marketing rules and systems and controls is sufficient to deliver a viable sector, and whether it is in any case the most efficient means of doing so.

Shiller (2015) argues that crowdfunding regulation concentrates too much on such traditional regulatory mechanisms when it would be more prudent to focus on reducing deceptive signals. Ben-Shahar and Schneider (2014) posit that retail investors lack the skill to make use of disclosure material for investing decisions. Many retail investors lack the time to conduct due diligence activities; others may be incentivised against this by investment caps enforced by regulators—for instance in the cases of the USA, Australia or Taiwan (Agrawal et al. 2013).

Given these facts, the most cost-effective way of improving market outcomes may be to focus on improving the quality of signalling in the market. There is precedent of such interventions in more mature and formal markets, with AIM in London incorporating requirements for some level of institutional involvement in funding rounds (Ibrahim 2015; Gabison 2016). New Zealand Stock Exchange's high growth segment for SMEs requires the separation of disclosure requirements into operational stages so that investors can more adequately monitor young companies (Pekmezovic and Gordon 2016).

A review of literature related to signalling in alternative markets suggests three potential approaches to strengthening signal quality:

1. Reining in oversubscribed offerings. In Bahrain, for example, any oversubscription to an offer must be refunded to investors. Similar regulations could constrain the ability of issuers to manipulate signals through setting artificially low funding targets which are then quickly surpassed signalling momentum to investors (Gabison 2016). By forcing funding consequences for the targets that an issuer presents, this approach better aligns the signal of how much has been raised with how much is actually needed.
2. Investor composition disclosures. Issuers might provide superior signals to investors by disclosing who has invested in their offering, as well as the investors' sophistication level and relationship to the issuer (Ibrahim *op cit*, Gabison, *op cit*). Jurisdictions such as Malaysia or New Zealand have in place restrictions or disclosure requirements for the actions of the issuer investing in their own securities through the platform; these could be expanded to include disclosing any investors with affiliations to the issuer and institutional investors. Italian crowdfunding regulation takes this a step further requiring at least five per cent of funding to be committed by

professional investors as a way of ensuring higher quality due diligence is undertaken on an offer (Gabison, op cit). A better understanding of who has invested will help other investors make decisions about the quality of due diligence already conducted.

3. Graduated investment offers. These link specific tranches of funding to corresponding milestones in business growth, and allow investors a 'cooling-off period' at the end of each milestone before the next tranche of their investment is committed (Ibrahim, op cit). Many regulatory frameworks incorporate one-off cooling-off periods for investors, but linking funding to milestones could greatly improve signal quality. South Korean regulators currently impose a post-funding-round share lock-up period for major investors while Taiwan has post-funding operational report requirements; combining these two and extending requirements to have different funding stages may further improve signals available to investors.

Even as good practices spread across the globe, complementary mechanisms such as the above are worth exploring further. They could, in time, prove to be major components of a second wave of global equity crowdfunding regulation.

Acknowledgements The authors would like to thank Yanzhe Li, PhD Candidate, Queen Mary, University of London, Centre for Commercial Law Studies, for her assistance and permission to reuse past research material.

References

- Aggarwal, R., and P. Stein. 2016. *The complex regulatory landscape for fintech: An uncertain future for small and medium-sized enterprise lending*. In World Economic Forum White Paper, 170816. http://www3.weforum.org/docs/WEF_The_Complex_Regulatory_Landscape_for_FinTech_290816.pdf.
- Aggarwal, Ajay, Christian Catalini, and Avi Goldfarb. 2013. *Some simple economics of crowdfunding*. NBER Working Paper Series, Working Paper 19133. <https://www.nber.org/papers/w19133.pdf>.
- Altfi. 2016. *Where are they now? 2016*, November. London. Altfi.com. https://www.altfi.com/downloads/Where_are_they_now_2016.pdf.
- Astebro, Thomas B., et al. 2018. Herding in equity crowdfunding. *Paris december 2018 finance meeting Eurofidai-AFFI*. 2019. https://sistemas.colmex.mx/Reportes/LACEALAMES/LACEA-LAMES2019_paper_319.pdf.

- Baeck, Peter, Liam Collins, and Bryan Zhang. 2014. *Understanding alternative finance*. NESTA and University of Cambridge. <https://media.nesta.org.uk/documents/understanding-alternative-finance-2014.pdf>.
- Ben-Shahar, Omri, and Carl E. Schneider. 2014. The futility of cost-benefit analysis in financial disclosure regulation. *The Journal of Legal Studies* 43: S2, Law & Economics Working Papers, 100. https://repository.law.umich.edu/law_econ_current/100.
- Blaseg, Daniel, Douglas Cumming, and Michael Koetter. 2020. Equity crowdfunding: High-quality or low-quality entrepreneurs?. *Entrepreneurship theory and practice*. <https://journals.sagepub.com/doi/abs/10.1177/1042258719899427>.
- BOP Consulting. 2017. *China UK equity crowdfunding regulatory and Market comparison study*, June, BOP Consulting. http://bop.co.uk/assets/others/China_UK_Crowdfunding_Regulatory_and_Market_Report_Final_24072017.pdf.
- Chen, Liang, Zihong Huang, and De Liu. 2016. Pure and hybrid crowds in crowdfunding markets. *Financial Innovation* 2: 19. <https://jfin-swufe.springeropen.com/articles/10.1186/s40854-016-0038-5>.
- Cumming, Douglas, Michele Meoli, and Silvio Vismara. 2019. Investors' choices between cash and voting rights: Evidence from dual-class equity crowdfunding. *Research Policy* 48: 8, October. <https://www.sciencedirect.com/science/article/abs/pii/S0048733319300228>.
- Cumming, Douglas J., and Sofia A. Johan. 2019. *Crowdfunding: Fundamental cases, facts, and insights*. Academic Press.
- European Bank for Reconstruction and Development & Clifford Chance. 2018. *Regulating investment and lending-based crowdfunding: Best practices*. 23 August 2018. <https://www.ebrd.com/news/publications/guides/best-practices-for-regulating-investment-and-lendingbased-crowdfunding.html>.
- ESMA. 2014. Opinion on investment-based crowdfunding. 18 December 2014. https://www.esma.europa.eu/sites/default/files/library/2015/11/2014-1378_opinion_on_investment-based_crowdfunding.pdf.
- Financial Conduct Authority. COBS. <https://www.handbook.fca.org.uk/handbook/COBS/10/2.html>.
- Financial Conduct Authority. 2018. *Loan-based ('peer-to-peer') and investment-based crowdfunding platforms: Feedback on our post-implementation review and proposed changes to the regulatory framework*. <https://www.fca.org.uk/publication/consultation/cp18-20.pdf>.
- Financial Stability Board. 2017. *Financial stability implications from FinTech: Supervisory and regulatory issues that merit authorities' attention*. The Financial Stability Board (FSB) Basel—Switzerland, 27 June. <https://www.fsb.org/wp-content/uploads/R270617.pdf>.
- Gabison, Garry A. 2016. The incentive problems with the all-or-nothing crowdfunding model. *Hastings Business Law Journal* 12: 489. https://repository.uchastings.edu/hastings_business_law_journal/vol12/iss3/3.
- Gregson et al. 2017. A review and simulation of business angel investment returns. *Venture Capital: An International Journal of Entrepreneurial Finance* 19 (4):

- 285–311. https://www.tandfonline.com/doi/full/10.1080/13691066.2017.1332546?casa_token=GasSEiLqc0cAAAAA:XR0sao_Szl6ersbZSwQPL5cYuFMgSVXSgdi9jv88wlcpYhO290K_FozKwcQQmcAmJCUWni86-wucyQ.
- Hornuf, Lars, Tobias Schilling, and Armin Schwienbacher. 2019. *Are equity crowd-funding investors active investors?*, CESifo Working Paper 7884. https://www.eco.nstor.eu/bitstream/10419/207275/1/cesifo1_wp7884.pdf.
- Hurst, Samantha. 2018. Saudi Arabia's capital market authority grants first two trial fintech licenses. *Crowdfund Insider*, 10 July. <https://www.crowdfundinsider.com/2018/07/136187-saudi-arabias-capital-market-authority-grants-first-two-trial-fin-tech-licenses/>.
- Ibrahim, Darian M. 2015. Equity crowdfunding: A market for lemons?. *Minnesota Law Review* 190. <https://scholarship.law.umn.edu/mlr/190>.
- Kisseleva, Katja, Aksel Mjos, and David T. Robinson. 2019. *The returns to early stage investment*, 28 March. https://www.glion.edu/wp-content/uploads/2019/07/9.Paper-KisselevaMjosRobinson_20190328.pdf.
- Lin, L. 2017. Managing the risks of equity crowdfunding: Lessons from China. *Journal of Corporate Law Studies* 17 (2): 327–366. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2908123.
- Lukkarinen, Anna, et al. 2016. Success drivers of online equity crowdfunding campaigns. *Decision Support Systems* 87: 26–38. <https://www.sciencedirect.com/science/article/pii/S0167923616300598>.
- Mamonov, Stanislav, and Ross Malaga. 2018. Success factors in title III equity crowdfunding in the United States. *Electronic Commerce Research and Applications* 27: 65–73. <https://doi.org/10.1016/j.elerap.2017.12.001>.
- Madsen, Joshua, and Jeff L. McMullin. 2018. *Economic consequences of risk and ability disclosures: Evidence From crowdfunding*. <https://www.smeal.psu.edu/accounting/research/economic-consequences-of-risk-and-ability-disclosures-evidence-from-crowdfunding.pdf>.
- Oxera. 2015. *Crowdfunding from an investor perspective: Final report prepared for the European commission financial services user group*. https://ec.europa.eu/info/sites/info/files/file_import/160503-study-crowdfunding-investor-perspective_en_0.pdf.
- Pekmezovic, Alma, and Gordon Walker. 2016. The global significance of crowdfunding: Solving the SME funding problem and democratizing access to capital. *William & Mary Business Law Review* 7: 347. <https://scholarship.law.wm.edu/wmblr/vol7/iss2/3>.
- Rau, P. Raghavendra. 2019. *Law, trust and the development of crowdfunding*. JBS Working Paper, University of Cambridge. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID3420705_code58471.pdf?abstractid=2989056&mirid=1.
- Rowan, Phillip, Bryan Zhang, Maragret Miller, Emmanuel Schizas, and Tania Ziegler. 2019. *Regulating alternative finance: Results from a global regulator survey*. World Bank and Cambridge Centre for Alternative Finance. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2019-11-ccaf-regulating-alternative-finance-report.pdf.

- Rowley, Jason D. 2019. Q4 2018 closes out a record year for the global VC market. *Crunchbase*, 7 January. <https://news.crunchbase.com/news/q4-2018-closes-out-a-record-year-for-the-global-vc-market/>.
- Samarah, Waleed E.A.R., and Saleh F.S.A. Alkhatib. 2019. Crowdfunding operations: Outreach factors in developing economies. *Journal of Public Affairs*. <https://doi.org/10.1002/pa.1988>.
- Schwartz, Andrew A. 2018. *Equity crowdfunding in New Zealand*. University of Oxford, Blog, 30 October. <https://www.law.ox.ac.uk/business-law-blog/blog/2018/10/equity-crowdfunding-new-zealand>.
- Shiller, Robert. 2015. Crowdfunding or crowdfphising *Project Syndicate*, November 15 New Haven. <https://www.project-syndicate.org/commentary/sec-crowdfunding-regulations-by-robert-j-shiller-2015-11?barrier=accesspaylog>.
- Signori, Andrea, and Silvio Vismara. 2016. *Returns on investments in equity crowdfunding*. SSRN 2765488. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765488.
- Walthoff-Borm, Xavier, Tom R. Vanacker, and Veroniek Collewaert. 2018a. Equity crowdfunding, shareholder structures, and firm performance. *Corporate Governance: An International Review* 26 (5): 314–330. <https://onlinelibrary.wiley.com/doi/10.1111/corg.12259>.
- Walthoff-Borm, Xavier, Tom R. Vanacker, and Veroniek Collewaert. 2018b. Equity crowdfunding: First resort or last resort?. *Journal of Business Venturing* 33 (4): 513–533. <https://www.sciencedirect.com/science/article/abs/pii/S0883902617304883>.
- White, Joshua T. 2016. Outcomes of investing in OTC stocks. *US securities & exchange commission DERA white paper*. https://www.sec.gov/dera/staff-papers/white-papers/16dec16_white_outcomes-of-investing-in-otc-stocks.html.
- Wiltbank, Robert E. 2009. *Siding with the Angels*. London. NESTA, May. https://www.nesta.org.uk/documents/51/siding_with_the_angels.pdf.
- World Bank and CCAF. (Forthcoming). *Global online alternative finance volume rankings and market volume index*.
- Xu, D., and M. Ge. 2017. Equity-based crowdfunding in China: Beginning with the first crowdfunding financing case. *Asian Journal of Law and Society* 4 (1): 81–107. <https://www.cambridge.org/core/journals/asian-journal-of-law-and-society/article/equitybased-crowdfunding-in-china-beginning-with-the-first-crowdfunding-financing-case/DFB441617948610F80B2E49D2D6BDCE6>.
- Zhang, B., T. Ziegler, K. Garvey, S. Ridler, J. Burton and N. Yerolemou. 2017. *Entrenching innovation—The 4th UK alternative finance industry report*. Cambridge Centre for Alternative Finance. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3084570.
- Zhang, B., T. Ziegler, L. Mammadova, D. Johanson, M. Gray, and N. Yerolemou. 2018. *The 5th UK alternative finance industry report*. Cambridge Centre for Alternative Finance. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2018-5th-uk-alternative-finance-industry-report.pdf.

Data-Driven Business Models



The Use of Psychometrics and Artificial Intelligence in Alternative Finance

Peter Romero and Stephen Fitz

1 New Sources and Use Cases of Data

Incumbents in the financial industry are tightly connected to political and corporate processes. However, current societal and technological advancements are transforming industries radically and at an unprecedented pace. Since the financial industry is heavily regulated and protected by the government, these transformations occur on two levels. First, incumbents are subject to the same technological and societal changes, and they deploy all new technologies to the extent they can. This deployment is however, limited by existing regulations, corporate strategy, and internal administrative inertia. Second, new opportunities for alternative finance companies emerge.

These new alternative finance firms can either include start-ups (often called FinTech) or existing corporations from other industries (often called TechFin), often from the computer industry. It is important to note that while TechFin corporations have the advantage of strong brand reputation and abundance of resources, they are limited by the same internal constraints

P. Romero (✉)

University of Cambridge, Cambridge, UK

e-mail: prr28@cam.ac.uk

S. Fitz

Keio University Global Research Institute, Tokyo, Japan

e-mail: stephenf@keio.jp

than incumbents in the financial industry. In contrast, FinTech start-ups are typically more agile flexible due to their flat hierarchies and the lack of established administration. However, most start-ups are also limited by the funding and strategy of their venture capitalists or investors, and face difficulties in recruiting talent. While in theory, FinTech and TechFin corporations compete for the same emerging opportunities, in practice, structural differences determine the kind of business they finally conduct.

Each of these three company types—existing incumbents, entrants from other industries, and start-ups—is at another stage of its transformation. Incumbents typically lag, but are powerful cornerstones of the financial industry. Corporations from other industries exploit their unique advantages in opening niches where they have competitive edges. Start-ups exploit different and more specialised niches. For example, while incumbent HSBC may have started using new algorithms for fraud detection and people analytics, Apple, a competitor from a different industry, expanded its Apple Pay solution to other countries and devices, and new start-up, TransferWise, developed a way to offer cheaper cross-border transfers and may experiment with social plug-ins. However, this situation is volatile. For example, Apple could decide to become a bank in some countries to offer customers cheque accounts, insurance products, or loans. This happened in several industries already, for example the UK supermarket chain Sainsbury's offers insurances, credit cards, loans, and saving accounts. German car maker BMW offers to its customers financing of cars, call money, credit cards, investment funds, and to its retail network merchandise and stock-keeping financing. TransferWise could start offering customer peer-to-peer lending. Challenger banks like Monzo in the UK already exist. While they are currently mobile-only and offer one service, the service is good enough to be a serious alternative to incumbents due to their potential for disruptive innovation (Wan et al. 2015). Beyond such start-ups, there are specialised services that deploy cutting edge technology like block-chain or occupy a very specialised niche like Patreon, a subscription service provider for content creators. While not being an immediate threat to incumbents, they are part of a growing infrastructure of alternative technology, which facilitates even further acceleration of the current developments. In this chapter, the authors identify three topics that broadly map the development of alternative finance.

1. Barriers to entry to incumbents
2. Evolving data trends in a changing world
3. Emerging FinTech and the pressure to adapt

1.1 Barriers to Entry to Incumbents

Banks and other financial institutions are strictly regulated with regard to which data they can use for improving or enabling their services, enforcing compliance, detecting fraud, conducting business intelligence or people analytics, informing marketing, and evaluating business partnerships. World Economic Forum publishes a comprehensive series of overviews on this topic together with leading financial institutions and universities, including one on FinTech (Trout et al. 2016).

Regardless, most banks started initiatives to make better use of their data by deploying emergent technologies like machine learning and big data analytics (Shah 2015). Thereby, they face a variety of organisational constraints. Their data is spread through different systems and lacks unique identifiers. These are common variables, with which data sets can be connected with each other. Often, they neither have modern tools or systems, nor the right talent at their disposal. Their senior leadership often lacks vision and therefore crucial investments are not made (Toesland 2017). A 2016 survey among 200 IT professionals and senior executives in the banking industry shows that more than 75% of senior executives believe that their companies are advanced in using big data, while this perception is shared by only 30% of junior managers and IT professionals. 54% of both groups identify regulations as the main challenge in the coming years, and 48% identify legacy systems as main challenge for IT. Only 32% identify economic pressure as a reason to conduct big data projects, while 45% state that their current systems cannot provide them with real-time insights (Shah 2015).

Compared to banks, Credit scoring agencies are limited in the quality and quantity of data they can use. For example, FICO uses the following data points to make their decisions:

- Payment history, including delinquencies
- Amounts owed, including amount of available credit
- Length of credit history
- Types of credits used, like revolving and instalment
- New Credit taken, including inquiries
- Number of recently opened accounts

FICO states that it will not use personal or demographic data like age, race, local address, marital status, income, or employment (Kaufman 2016; McFadden 2010). German SCHUFA makes its decision based on the following data (Vergleich.de 2018):

- Account openings
- Mobile contracts
- Leasing contracts
- Monthly payments
- Frequency of loan applications
- Duration of living at same address
- Requests for internet shopping
- List of debtors
- Late payments, for example of mobile bills
- Affidavit or arrest warrant to enforce affidavit

SCHUFA officially states that it will not record nor use socio-demographic information for their scoring. This could encompass such informations like the nationality of individuals, their age, disabilities, religion, or ethnicity (SCHUFA 2018b). In 99.7% of all cases it makes no use of geographical data like the quality living areas, the percentage of families, or consumerism in that area. They state that they only use such data, when there is not enough alternative data about individuals available. They explicitly state as well, that they do not use data from social networks for creating their scores (SCHUFA 2018c).

However, SCHUFA is an opaque organisation, which influences many aspects of life in Germany and faces a lot of criticism. At the moment, an initiative of data scientists and statisticians called 'Open SCHUFA' asks people to donate their credit reports so they can reverse-engineer the SCHUFA algorithm (OpenSchufa 2018). This, in turn, provokes criticism from SCHUFA, which argues that this would help people with a low credit ranking to trick the system, and that the algorithm is their IP (SCHUFA 2018). The initiative is still ongoing and reports first results, which contradict the official statement of SCHUFA about the data used to create the score. They conclude indications for a broken procedure, lack of trustworthiness, discrimination through the score, and biased algorithms. Persons with no negative characteristics are rated high risk, reports have missing data, factors like age, sex, and frequency of moves may influence the rating, and it appears as if there are some outdated versions of the algorithm still in use. Up to today, 'Open SCHUFA' collected just 3,000 data donations, therefore these first results could change still (OpenSchufa 2018).

Furthermore, these results cannot be verified since the SCHUFA algorithm is considered to be a trade secret (SCHUFA 2018a) and therefore not available to the public. In 2012, SCHUFA was involved in a data research project

called 'SCHUFALab@HPI', which created a public outcry. They cooperated with the renowned Hasso-Plattner Institute to use social media data for augmenting their credit ranking system. Their goal was to incorporate the following data points into their algorithms:

- Contact network of a person
- Analysis of text to create a 'current range of opinions of a person'
- Detect addresses and address changes through a person's network
- Creating an unspecified 'pool' of data for existing or future products and services

Beyond securing their 'quality leadership', their goal was as well to conduct scientific research on 'countless myths and assumptions' about the 'information source internet' (Wilkens 2012). At that time, Facebook, Instagram, Snapchat, and mobile media were not permeating society as strongly as they do today, and before Edward Snowden, Cambridge Analytica, and the Facebook data breaches, the public was less sensitised to privacy topics. Also, due to the prevalent criticism of SCHUFA and the low German willingness to share data (Guenole et al. 2015), this project was halted soon.

In summary, incumbents of the financial industry are constrained by regulations, administration, protectionism, lack of vision, public scrutiny, low levels of trust, dysfunctional communication, scattered and incompatible data, outdated systems, and lack of technical talent. Though they display willingness to deploy the latest technology, they do not make use of the full potential of cutting-edge technologies like machine learning or big data analytics. The authors conclude that while incumbents tackle technological challenges, they would benefit from focussing more on societal changes, as well. Future research should therefore focus on how incumbents can align organisational transformation with societal changes, and how alternative finance companies could exploit these weaknesses through direct competition or through offering incremental B2B services to inert incumbents.

1.2 Evolving Data Trends in a Changing World

Today, a multitude of data points are used for analysing people, making decisions on them, providing assistance, and nudging them towards a desired direction. The reader will find in the appendix an exemplary list of data points that are collected through today's technology on a regular basis, and what data scientists and psychometricians can deduce from these. Compared to that, the 'SCHUFALab@HPI' project is rather non-intrusive. The amount

of available data grows explosively, and in just two years, 90% of the world-wide collected data was produced. Between 2015 and 2017, there was a 42% increase of people using the internet, adding up to today's 3.8 billion internet users (Schultz 2017).

By 2025, the global amount of collected data is forecasted to be 175 zettabyte, which is 175 trillion gigabytes. This data will be produced by large data centres, public and private clouds, through enterprise servers and cell towers, and through countless endpoints like cars, mobile phones, robots, computers, Internet of Things (IoT) devices, wearable technology, and augmented reality devices. It will be produced, manipulated, and processed almost in real-time, which means that 150 billion internet-enabled devices will permanently communicate with data centres and with each other through the 5G network that is built up at the moment. The percentage of real-time data from the entire data-sphere will have doubled from 15% in 2017 to 30% in 2025. This trend is supported by the change of storage location from local to cloud servers, whereby the overall cloud usage is going to double between 2017 and 2025. The biggest chunk will be in public or private cloud services then (Reinsel et al. 2018). Internet usage will grow in the developed and in the developing world. However, the developing world displays an uninterrupted increase, whereas the developed world shows signs of saturation (Wikipedia 2018). Therefore, more data will be on the move, which demands faster connection speeds, more central and cheaper storage, and the necessity to analyse it pre-emptively in a decentralised mash of devices. At that time, it might be even more important to decide what to forget than what to save.

The authors identify several trends of data processing in organisations that emerged recently. First, data processing and insight generation becomes more automated. Where teams of statisticians and econometricians worked before for weeks on a problem, one data scientist equipped with tools like Alteryx, SparksBeyond, or IBM Watson can now do the same task in hours. That means, that bottom-up, data-driven approaches grow in importance over top-down, theory-driven approaches at one point in time. Second, and counter-intuitively, focus on giving meaning to existing data could have strong benefits over collecting ever more of it. Third, mobile and social data is the centre of attention of most organisations, based on its ubiquity. This encompasses personalisation to individuals, interactive engagement of customers on various channels, as well as activating broader social networks through influencers. Last, it is more important to gather a holistic understanding of customers' lives than knowing only one aspect of them in

greater detail. This can be achieved by sharing data or insights with other organisations and by eliminating own data and cooperation silos (IBM 2017).

These trends correspond with the general recommendations on ‘The Appropriate Use of Customer Data in Financial Services’, a study from World Economic Forum in cooperation with Oliver Wyman (World Economic Forum 2018). They explore the implications of the fourth industrial revolution on the financial system and its actors, and conclude the need for global guiding principles for finance companies to prevent a fragmentation of frameworks and regulations:

1. *Control*—necessity to be aware of what is done with the data and why, and whether consent is given.
2. *Security*—responsibility for security and data protection
3. *Personalisation*—freedom to create personal profiles of customers on individual level
4. *Advanced Analytics*—ability to test, validate, and explain models to customers
5. *Portability*—access to data for customers and third parties

They propose that public regulations and organisational best practices should be merged into a collective approach, and that incentives should be used to accelerate adaptation. However, the authors believe that organisational constraints of incumbents will make a rapid adaptation difficult, and that small alternative finance companies and FinTech start-ups are best positioned to exploit these changes, adapt in the necessary speed, and that they will ultimately change the financial industry through disruptive innovation. The authors conclude that future research should identify optimal strategic positions that yield the highest potential for disruption. A special focus should be put on societal changes and their implications for organisational leadership and desired employee behaviour.

These societal and political changes facilitate and normalise enhanced collection and connection of data. For example, Estonia’s e-residency allows every person of any nationality to open bank accounts, found companies, access international payment providers, and get taxed there. Its holders share all relevant data via the internet with Estonian public services, which store and process it centrally. Her national electronic ID system is necessary for all aspects of life like identification, travel, national health, signatures, medical tracks, and prescriptions. There is a political discussion whether social media and internet accounts should be connected to that ID system, as well. This would make the online behaviour of every citizen uniquely identifiable, but

could be used to fight the spread of fake news and attacks from computer criminals. This approach is comparable to a Russian law, which demands all online services to use mobile phone numbers as a login. Through this, identity theft can be massively reduced and online orders and postings in social networks are uniquely identifiable. Similar proposals are being discussed in the European parliament, as well.

An increasing number of countries demand biometric markers like photos and fingerprints in their passports. India leads this development with the world's largest biometric database, called Aadhaar. It incorporates all fingerprints, iris scans, and biometric facial photographs of her citizens. Aadhaar officials would like to connect it with more data like health records, credit ratings, electronic signatures, criminal files, and online behaviour. An Aadhaar number is mandatory for tax files, returns, and to claim welfare benefits. Banks already offer cash withdrawal and paperless banking through it, which provides new opportunities to the under-banked and poor people. Once made accessible to third parties, it could enable new FinTech start-ups or companies like Facebook to uniquely identify individuals and offer online payments through the Aadhaar network (Arora 2017; DNA India 2018). After a law suit with privacy advocates, corporations are not allowed yet to ask for the Aadhaar number, however discussions around that are still ongoing (Safi 2018).

China created a 'Social Credit Score' for her citizens that is operated by Baidu, Alibaba, and Tencent, who also operate most Chinese social networks and the duopoly of mobile payment providers, AliPay and WeChat. This score is visible to everyone, encompasses all data used by credit score providers, and also measures the political compliance of individuals and their friends, for example through social media, shopping behaviour, police data, surveillance cameras, and administrative data. Financial services are at the core of that score, since it is almost impossible to pay for anything in China without WeChat or AliPay. Therefore, they can interlink online and offline behaviour with daily payment processes to compute the social credit score, but also to tailor bespoke advertisement and trigger desired behaviours (Botsman 2017).

Many of the technologies of the Chinese 'Social Credit Score' system have Western counterparts. These are not yet as connected with each other as their Chinese counterparts, nor is the data officially integrated in one central governmental database. However, after the PRISM scandal, the reader may assume an unofficial integration through the NSA system XKeyscore (Nolan 2013). Amazon offers a system called 'Rekognition' (*sic!*) for connecting public or private cameras, which recognises faces, matches them against a

database of people of interest, and tracks the movements of these people, even in crowded places. The same system can be used for analysing content that people load up to internet servers, for example, to identify adult or criminal material (Amazon 2018). Microsoft's Windows 10 receives strong criticism for the amount of data it collects from its users, and the tight integration with its cloud services. They analyse the content users store on their cloud services to offer additional features and to enforce their terms of service. For example, they provide automated transcripts of videos for Office 365 customers (Finnegan 2018) and scan for illegal content and such content that is not in line with their Code of Conduct, for example, partial or full nudity in photographs, art, or drawings (Microsoft 2018a). To analyse pictures, they use their software PhotoDNA, which is as well in use with law enforcement, Twitter, Facebook, and Google. The latter just forbid pornographic content on their cloud services. Therefore recently, several sex workers lost business since their professional content was erased or locked down from their Google-provided cloud storage (Cole 2018).

Beyond that, Google analyses most of the data of their users to create consumer profiles for advertisers. Recently, they considered offering a localised Chinese version of their search engine that is aligned with the Chinese government's strict regulations and surveillance directives for population control. Google also developed its own visions for behavioural modification in an 'internal thought experiment', called the selfish ledger. This is an AI-driven personal assistant and decision helper that not only uses all of the users' data to understand them better and nudges them into certain desired behaviours and, of course, to provide even more data, but also aligns the behaviours it recommends with the values of Google as an organisation. Since Google offers a mobile payment solution and has increasing knowledge about the offline purchasing habits of its users, this approach would encompass financial decisions, as well (Savov 2018). Also Facebook is interested in behavioural modification. They conducted a series of experiments on their users to change their mood and voting behaviour by manipulating their news feeds (Bond et al. 2012a; Kramer et al. 2014), which had influence on a broader societal level (Stoycheff 2016). However, it could also be used to improve the effect of advertisement and increase the tendency of people to buy certain goods.

Access to user data and control over the information displayed to them would yield unprecedented control over the private lives of citizens. Current legislations could enable exactly that. Should the EU Copyright Directive Articles 11 and 13 be passed, providers will have to filter the data that is being uploaded in real-time. While enforcing copyrights is a legit demand,

it could also be used to prevent fake news or hatred to be spread. Access to data helps law enforcement already in investigations and to prevent crimes. Paypal co-founder Peter Thiel founded a major company for predictive policing, Palantir. Their software Gotham and Metropolis allows to collect and analyse structured and unstructured data in real-time, display trends, recognise patterns, and predict future scenarios. It is used for policing, warfare, governmental agencies, and private companies. But it is also used in the finance industry. Banks and hedge funds use it for example to identify insider threat and to predict fraud (Fishman and Hassebrock 2010). Based on the trends of data requests from law enforcement and governments, and the ongoing cross-pollination of civil and official applications, the authors conclude that alternative finance companies will benefit from segmenting their markets. Thereby, the transfer of existing solutions to new applications should be in focus. Also, products and services should be explored that connect data from different domains to better understand, predict, enhance, and change human behaviour.

By connecting online behavioural patterns with each other and with offline behaviours, companies and governments can create holistic psychological profiles of individuals that cover likely behaviours, values, personality, ability, hopes, needs, preferences, and fears. These capabilities are strengthened by a growing number of Internet Of Things (IoT) devices, which produce an increasing quality and quantity of behavioural data, and which become increasingly personal and 'intimate', closer to the brain. From a Desktop to a Laptop, to a smart phone, to a smart watch, to smart glasses and lenses, to direct brain communication that many companies are working on, for example Facebook, DARPA, and Elon Musks's Neuralink. This miniaturisation shifts the centre of information processing from individual devices to meshed device networks, federated clouds of ambient computing that collect, store, pre-process, and pre-evaluate data autonomously. Examples for this behaviour-oriented, ambient, and closer-to-brain development are:

- Geospatial identification of customers through Google that enables payment through their device location.
- Apple's and Google's Bluetooth beacons that identify customers, update their advertisement profiles in real-time, and present information on their phones.
- Apple's, Samsung's, and Google's mobile payment systems in smart phones and smart watches with biometric identification.
- Augmented reality applications for displaying additional product information in physical stores, and collecting more data about the customers.

- Wifi ID catchers that recognise mobile devices in malls, subways, and buildings to create movement profiles of customers.
- Humanyze's employee batches that register who speaks with whom where in the office, analyse the sentiment of their communication, and generally maps who is where in the office.
- Built-in Office 365 People Analytics capabilities, which was one of the reasons Microsoft acquired Volometrix.
- Facial recognition of customers in stores, that directly connect to their Amazon or Facebook accounts, and enable contact-less payment through RFID chips on the products.
- Lower insurance rates in exchange for location data in cars and fitness data from fitness devices.
- Ultrasound signals in television programs that interact with mobile apps and detect which customers are watching which television series.
- Smart contact lenses, attachable e-tattoos, and radio-frequency identification (RFID) implants that allow to make payments or open doors.

New use cases that are already under development that become increasingly ambient and personal. For example, the Microsoft Shopping Assistant uses data from the smart phone, computers, and web browser of users. It remembers the products they are interested in, tracks prices, and sends notifications through various channels like email or Facebook Messenger (Microsoft 2018b). In the near future, it could be sending information to an implant that directly communicates with customers' brains, and aligns its recommendation with their purchasing history, friends network, financial situation, plans, emotions, location, health situation, and the overall political strategy of the company or country that the customers are in (Berens 2017). Implantable technology could provide social nudges by informing its wearers about other people's behaviour for shaping norms and behavioural expectations (van der Linden 2018). They may as well directly manipulate behaviour (Schiff et al. 2007) through brain stimulation or synthetic rewards (Garris et al. 1999). This could be used to shape general pro-societal behaviour and the use cases and boundaries for policymaking over such devices are already discussed (Rodotá and Capurro 2005).

That may sound dystopian, but the idea of machines doing human labour was dystopian for the Luddite in the nineteenth century, as well. The first industrial revolution laid the foundation for our wealth and the subsequent societal changes created stable democracies. Alternative finance companies should be prepared for rapidly changing societal norms, omnipresence of computing and surveillance, implantable technology, and a growing trend

towards more collectivist societies with strong and enforceable behavioural norms. Further research should be done on understanding how alternative finance integrates with other applications and services that are connected in the Internet of Things (IoT). Using the logic from system theory (Willke 2000), the authors predict that possible applications and business opportunities will emerge at the interfaces between and within elements of this IoT, where financial services are not only used to regulate the system but also to exchange information. A first supportive indication are alternative applications of block-chain technology like Ethereum, which can be used not only for payments, but also for identification, operational optimisation, betting, sourcing and pricing, and communication.

1.3 Emerging FinTech and The Pressure to Adapt

Growing societal and technological connectivity enables new financial services or products that change the expectations and needs of customers. While expectations define what customers expect to find with products and services based on their habits, needs are much more fundamental and their absence may lead to adverse outcome (Maslow 1943). Customers feel disappointed if their expectations are not met, however they feel deprived if their needs are not met, and they will either complain or look elsewhere to satisfy them. Expectations can become needs over time, depending on how much customers take them for granted and to which degree they are necessary to self-sustain or participate socially. For example, internet and WiFi in hotels and meeting centres was an expectation a couple of years ago, but for many, it is a need today.

TechFin companies develop ecosystems of products that are interconnected, easy to use, predictable, personalised, addictive, and designed to engage users through positive experience (Capgemini and LinkedIn 2018). Their offers of financial products and services are an integral part of their ecosystems, with the same predictable usage, feedback, and satisfaction patterns. The development towards ambient and personal computing is based on these ecosystems, and FinTech companies offer products and services in it, as well. Through the connection of these ecosystems with the everyday lives of customers, these demand the same experience from the incumbents. This imposes pressure on them, since they are constrained by their own organisational inertia, overheads, and strict regulations. In this chapter, the authors explore therefore both phenomena, and discuss their ramifications for research and practice:

1. Emerging challengers shape customer expectations
2. Evolving customer needs pressure incumbents to adapt

1.3.1 Emerging Challengers Shape Customer Expectations

Capgemini and LinkedIn (2018) find that compared with incumbents, Alternative Finance companies offer 'superior customer experience' (p. 18) through a 'customer-centric approach' (p. 27). Both are enabled by emerging technologies, agile development, and customer-centric design principles. However, implementation is driven by competition and rising customer expectations. The core emerging technologies are mostly developed by TechFin organisations: Operational Excellence (robotic process automation, chat-bots, distributed ledger technology), Infrastructure Enablers (platforms, cloud applications, automation programming interfaces—API), Front-End Interface (intuitive user interfaces, gamification, augmented and virtual reality), and Data-Focused Technologies (analytics, artificial intelligence, machine learning, sensor-based technologies, biometrics). FinTech companies mostly focus on niche applications and use the API of TechFin ecosystems to offer their products and services. This increases the agility and innovativeness of these ecosystems, makes them more relevant to customers, and contributes to improving their experience with regard to convenience, personalisation of experience, and functional pain points like responsiveness, turnaround time, and number of interactions.

However, better ecosystems also raise customer expectations through better convenience, intuitive interaction, more relevant functionality, and proactive insights based on user data and interaction. Alternative Finance companies have interlinked design and innovation processes that are centred around fulfilling the expectations of their customers and focus on simple user interfaces, as well as quick, convenient, and seamless interaction. Thereby, mostly only one functionality is addressed, which is either underserved by incumbents or which customers experience as of high friction. To increase relevance for customers, applications are personalised and driven by insights from existing data. Since development is agile, feedback can be incorporated and expectations of customers can be addressed quickly. Through embedding in the ecosystem, delivery of services is seamless and aligned with customer expectations. The underlying goal is that of a single customer journey, not that of individual touch points. This is realised by empathising with customers, simplifying their experiences, and aligning with their expectations and goals (Capgemini and LinkedIn 2018).

These ecosystems get increasingly useful with the number of their customers' areas of life they permeate, and the more holistic their situation is represented. Field theory from psychology helps understanding the dynamics of holistic personal situations. It integrates '...divergent physiological, psychological, and sociological facts on the basis of their interdependence. To explain social behavior it is necessary to represent the structure of the total situation and the distribution of forces in it' (p. 868). Thereby, Behaviour (*Be*) can be described as the function of a person (*P*) in a specific contextual environment (*E*), which equals to the function of a life-space (*LSp*).

$$Be = f(P, E) = f(LSp)$$

This life-space encompasses the person, other persons in their social network, their goals and resources, interaction, and behaviours, adjacent social groups, potential organisational settings that indicate hierarchy or greater influence, and the greater societal context. Persons 'unlock' different areas of their life-spaces through exploration and maturation. These areas can be geographical (e.g., entering school), social (e.g., changing societal classes), or behavioural (e.g., forming new habits) in nature. Thereby, desirable, similar, and wanted areas are explored, and undesirable areas are left out—for example new and good areas of cities, or prestigious social groups. Over time, persons learn to better differentiate between wishes and expectations. However, any time they decide to explore new areas, they are *in statu nascendi*; transitioning from one status to the other, and during this transition, they are in a formative state until they become 'established' in their new positions and relationships again. These areas of life can be independent or interdependent from each other, and the connection is made by the persons 'unlocking' them (Lewin 1939). This connection-making and the formative status of persons is an excellent opportunity to position products or services, especially when they offer familiarity, since it provides persons with perceived security. Therefore, companies have to know as much as possible about the specific situations of their customers, to change their behaviours, and program them to fulfilling emergent needs by using the products of that respective company. In such situations, familiar or similar persons have a stronger influence on customers than other people. For example, Facebook offers advertisement that displays faces of the core social network of persons, or such of influencers in their wider network.

During such a journey towards a desired state, driving forces are constrained by restraining forces inside and outside the person. Driving forces could be for example the general need for change due to stagnation or external events, existing conflicts, maturing, personal interests, demand at other places, uncertainty, competition, values, technological change, politics,

or globalisation. Restraints could be the structural inertia, sense of insecurity, anxiety, perceived stress through endangered adaptation (Selye 1955), bias by self-centricity, or habits (Lewin 1943). Persons will initiate change when they are forced to, dissatisfied with the current situation, or when they strive for a better situation. Their awareness of the forces that influence their situations and their willingness to act upon it grows with decreasing psychological and geographic distance of the factors driving change to their individual life-spaces. Also the strength of the driving and restraining forces on behaviours depends on how close the origins of these forces are to persons. The influence of the proximity of these origins to the person increases faster with avoiding than with approaching behaviour Miller (1951, 1944). The more difficult the tasks behind these behaviours are, the more difficult it is for persons to display correct behaviour, and the more adverse reactions are induced. However, simple tasks almost instantaneously trigger correct behaviour (Spence and Taylor 1951). Thereby, persons with lower degrees of fear produce less mistakes with difficult tasks than persons with high degrees of fear. However, during easy tasks, persons with higher degrees of fear produce less mistakes (Spence et al. 1956). Applied to Alternative Finance, the easy, ambient, known, and intuitive products of FinTech and TechFin companies outsmart those of the incumbents by lengths. Finally, the strength of the behaviour (E) is determined by an interaction of the strength of the drive for an result (D), the habit (H), the intensity of the dynamism (V), and the incentive in terms of personal value of the goal (I). The strength of the drive is mainly determined by the deprivation of the incentive, that of the habit is determined by the amount of successful behaviours upfront, that of the dynamism by the intensity of an individual stimulus, and that of the incentive by the value a person and the context attribute towards it (Hull 1952). The easier, more joyful, more personally relevant, and more habituated the use of a specific product from a TechFin ecosystem is, the less attractive those of incumbents become, given both fulfill the same utility and enjoy comparable levels of trust.

$$E = D \times H \times V \times I$$

Several drives can multiply and increase the probability of displaying behaviours, for example hunger and the need for social exchange may lead to accepting restaurant loyalty cards. However, also emotional states like moods, cognitive effects like distinction of objects, or situational effects like bad weather are factors that lead towards behaviours (Forgas et al. 2009). In summary, cognitive, emotional, biological, contextual, and social factors determine whether a certain behaviour is displayed or not. One of the

strongest factors for displaying behaviour is through social influence. As social beings, humans strive to be accepted by other humans and groups. The more prestige membership of a group generates, the stronger the willingness of persons to adapt. Once accepted, in-group behaviour will be internalised to a degree that it becomes indistinguishable from individual behaviour (Mason et al. 2007). However, group membership also induces competitive behaviour. Group members strive to attain the same goals as other group members in a mimetic desire (Girard and Williams 1996). If goals cannot be attained, and if behaviours cannot be displayed due to various constraints, compensatory behaviours could be displayed. For example, some persons may want an iPhone but can only afford Android devices. In order to maintain their positive self-image, they will start looking for a positive narrative that explains why they made the right choice, for example, the more open architecture of that mobile operating system. The selective market introduction of iPhones with their limited availability made use of that effect. The low availability and high prestige increased the demand, especially within social groups where first early adopters already had one.

Since mature persons are more conscious of their needs and expectations and are more settled in their life-spaces (Lewin 1939), younger people tend to be more often early adopters. For example, younger users are familiar with the ecosystem of products and services of TechFin companies, and they are happy to share personal data with them in various areas of their lives. Trust is developed through adaptation in their social circles and life-spaces, reliability, social confirmation, and through a series of positive experiences. And, since these products and services are simple and joyful to use, and facilitate connecting independent areas of life-spaces, strong user-habits are built-up within the ecosystems of TechFin companies and their FinTech partners. Therefore, users of these ecosystems trust TechFin companies to a degree that is close to the trust in incumbents, however FinTech companies need visible cooperations with incumbents or TechFin companies to gather the same level of trust from customers. They will profit from explaining why customers need these services, provide them with transparent information, offer assistance, and build up personal relationships. Through this, users perceive sharing data and implementing these ecosystems in ever more areas of their life-spaces as natural, satisfying, and convenient. Over time, this will build up trust, especially when users have the feeling to be in control of their data, experience transparency, are educated in a simple way, and know about the data protection measures of Alternative Finance companies (Capgemini and LinkedIn 2018).

Bain & Company finds that among Google, Apple, Facebook, and Amazon (GAFA) and other Alternative Finance organisations, customers would trust Amazon and Paypal the most with their personal banking du (Toit et al., 2017). Both companies permeate and connect various areas of their users' life-spaces; while PayPal is the *de facto* standard for online payments, Amazon offers globally a broad range of goods and services, including groceries, books, clothing, tablets, phones, movie and music streaming services, one of the largest commercial cloud infrastructures, streaming, social media, food markets, plumbing services, transportation services, payment services, and delivery services that can remotely enter customers' houses if they so wish. With the Echo device, they spearheaded the category of 'smart speakers'; always-listening internet-connected speakers that connect customers to an artificially intelligent digital assistant, 'Alexa', within their homes, household appliances, and soon within their offices, computers, smart phones, and cars. This makes Amazon one of the TechFin companies with the highest degree of penetration of all areas within people's life-spaces. In return, customers trust them with their data, which allows Amazon to offer bespoke services and products. Their strategy of low margins and permanent exploration of new services and product groups allows them to quickly react to customer expectations and needs. This imposes strong pressure on the players and competitors in the markets, segments, and sectors they enter, and there is no reason to assume that this pattern would not repeat in the banking sector as well, if they chose to enter it.

However, there are indications, that they are about to do just that. Bain calls it 'Banking's Amazon Moment' (p. 1), and describes it as a 'game changer' (p. 1) and an emerging existential threat for the banking sector and overall finance industry. Amazon currently has a co-branded credit card offered by Chase, customers can deposit cash on 'Amazon Cash' through a network of 'retail locations', and in 2017, they gave more than \$1 billion loans to small businesses selling through their platform. It is reported that they are in discussions with several incumbents, including JPMorgan Chase & Co. to offer a co-branded cheque account that is debit-based and mobile-centric. Though debit accounts are not profitable, this move would fit in Amazon's overall strategy to extend into markets and areas of life-spaces of customers first, and monetise later. Also, the incumbent would be responsible for managing the regulatory constraints, while Amazon would focus on the user-facing front-end. And since Amazon could do business with itself, they would save the 2% interchange fee for most transactions within their organisation and its third-party merchants, which sums up to more than a quarter of \$1 billion annually in the US alone. Once established, and

following their strategy and the logic outlined above, they are expected to move into other products and areas like lending, mortgages, casualty and life insurance, and wealth management. This means, they would also impose a threat to the insurance industry. Last, they don't have to maintain a network of branches, but could leverage on their Alexa/ Echo technology (du Toit and Cheris 2018). In return, the incumbent could obtain a distribution advantage through access to all relevant Amazon customers. And, both companies could increase their understanding about the connected life-spaces of their customers through merging available data. Amazon could get all account information and would be able to understand the holistic situation of their customers and their offline purchasing habits. The incumbent could get access to Amazon's understanding of the consumption and living habits of their customers and their online purchasing habits. Merged, this would yield in an unprecedented quality and quantity of data and potential for personalisation, which could potentially disrupt the financial industry (du Toit and Cheris 2018).

Further indications towards Amazon's interest in entering the financial industry is that it obtained an e-money licence from Ireland through its affiliate Amazon Payments Europe. This is not a full banking licence, but the option to create, store, manage, and transfer digital money. This allows its customers to do payments, however it does not replace cheque accounts and does not allow Amazon to offer credit cards. Facebook obtained such a licence from Luxembourg, and Google in Lithuania through its affiliate Google Payment. Google plans to offer users of its digital purse 'Google Wallet' an additional payment option to credit cards and PayPal. While Apple offers such a solution with its Apple Pay feature on their devices, the 1.5 billion customers with Google accounts impose a stronger threat to the banking industry due to scale. For most banks, such offers from GAFA mean loss in fees and—most importantly—access to customer data. With the e-money licence, GAFA can now circumvent the banks completely. Until January 2019, Lithuania issued 39 of these licences and the UK 128—mostly to TechFin and FinTech companies like TransferWise and Wirex. This indicates a growing network of alternatives for customers, represented by the ecosystems of TechFin and the embedded FinTech companies. Beyond better products and cheaper service, customers lost trust in many incumbents due to their unethical behaviour during the 2008 financial crisis (Mühlbauer 2019).

This development is more progressed in Asia. Chinese affiliate of Alibaba, MYbank, uses customer data from their parent company to approve loans instantly. Its other affiliate AliPay handles \$1.7 trillion on annual transactions, which is five times the amount PayPal handles. Rakuten started like Amazon

with e-commerce, and is today Japan's largest internet-based bank and third largest credit card company; their financial services division now generates 40% of the group revenue. Within the West though, Amazon is best positioned to disrupt the financial industry. They already hold personal data and credit cards of their customers, developed a strong relationship with them, offer a consistent and simple experience, integrated their products in nearly every area of their customers' life-spaces, and are integrated on computers, smart phones, tablets, and TV. They didn't have a security breach. And they are so diversified that they generate revenue with other business divisions and can operate on very low margins over an extended period of time (du Toit and Cheris 2018). Most importantly, they satisfy evolving customer needs.

In the UK, only 45% of bank customers agree that the homepage lets them do what they want, and only 25% of the bank customers think that the mobile app users think so. 45% think that the homepage is easy to use, and 34% think so about the mobile app. In the US, 27% of the bank customers use voice assistant on their phones and 18% already own a smart speaker. Throughout US, UK, and Australia, 5-6% of the bank customers already use a digital assistant for banking, however $\frac{1}{5}$ to $\frac{1}{4}$ of all customers would be open to try it in the future. Once in an ecosystem, users display a 'virtuous circle of higher adoption and digital first habits, contributing to greater loyalty' (p. 4). Customers start with digital interaction (87% with digital leaders of the industry, 59% with digital laggards of it), then they are successful (90% with leaders and 87% with laggards), which leads to an higher adoption rate (95% leaders, and 78% laggards). This positive experience contributes to a higher Net Promoter Score[®] (76% leaders, 37% laggards). However, among the banks, being a digital leader or a laggard not only makes a difference in new customer acquisition, but also in operating costs, especially in low-margin regions. Customers of the laggards act as detractors and resolving disputes is twice as expensive for them as for the leaders. This digital divide is based on affinity, age, and market. Younger people prefer digital solutions, and so do customers in markets like China, India, Mexico, and Brazil, where banks underserve. Asian customers are most active in using mobile payment apps, and developing countries favour digital purchases (du Toit et al. 2017).

That means that it makes sense to migrate routine interactions to digital channels, and assist customers on their journey. In the UK, this increased the Net Promoter Score[®] for routine interactions as a proxy of customer loyalty by 75% and decreased cost per interaction by 95%. However, customers are not eager to use AI only for complex interactions, but prefer mixed channels of AI with online chats and video conferencing with real humans. However, they are annoyed when they start digital for a routine interaction and have to

switch to a human operator for resolution. This effect is stronger for online interactions (drop Net Promoter Score[®] by 36 points from 45 to 9) than with mobile interactions (drop by 14 points from 50 to 36). In general, customers prefer human operators with emotional episodes that yield potential for a *catharsis* or moments of truth, for example when handling complaints or ‘demanding outstanding customer service’ (p. 5). Older US customers prefer in general digital assistants for routine tasks, and employees for emotional tasks, however younger US customers use for emotional episodes more often digital assistants. Finally half of the US consumers would get a credit card from a TechFin organisation, if these would offer one (du Toit et al. 2017). The authors see full alignment with psychological theory and human reaction on changes towards an increasingly digitised world, and expect these numbers to rise within the next five years, when digital natives from Generation Z will display increased participation in economy, and the Baby boomers and first members of Generation X retire. Future research would benefit from understanding change and customer migration strategies better with regard to cultural and age differences. Also, research should be conducted on ethical boundaries, and how data protection can be guaranteed without compromising value for the customer. Finally research and adjacent new product development should be conducted that explores how the addition of new areas of the life-spaces of customers increase their satisfaction and willingness to stay within the ecosystem. However, this should be aligned with the debate on ethics and data protection to create the foundation for a sustainable growth.

1.3.2 Evolving Customer Needs Pressure Incumbents to Adapt

TechFin companies can generate actionable insights about customer expectations from volumes of data they aggregate about them. They have a large customer base and strong cash reserves that allow scaled global operations. They develop and deploy emerging technologies and novel design principles that define optimal customer experience. While in Asia, their influence on the finance industry is strongest, e.g., Alibaba’s Ant Financial, Amazon is the best set-up to disrupt the West, since they can provide a consistent experience throughout most life-spaces, focus their business around innovation and diversification, and have an effective infrastructure. Most importantly, they are agile, have a high-performance culture, and a strong leadership with a complete vision (Capgemini and LinkedIn 2018). This nascent potential to disrupt the financial industry implies the question of how incumbents should adjust their strategy, and which opportunities for FinTech companies may

emerge. Capgemini and LinkedIn (2018) conclude that incumbents should cooperate with FinTech companies. These bring agility, focus on customer experience, ability to quickly release products that meet customer expectations, and absence of legacy systems. Incumbents bring a large number of customers, a known and mostly trusted brand, an existing distribution infrastructure, regulatory experience, and sufficient capital for scaling operations. Also, incumbents are increasingly willing to cooperate since they lack the talent and expertise for creating specialised services, as the FinTech companies can do. While some still build their own solutions, most look for external partners, for example in innovation incubators, business accelerators, hackathons, or at venture funds. However, the biggest challenge for incumbents and FinTech companies is finding the right partner. While until now, 7,500 FinTech companies were founded, and raised over \$109.8 billion, most of them fail due to wrong product-market fits, products that are easy to copy by competition and other missing isolating mechanisms, high costs for scaling operations and occupy market niches, and the inability to find the right partners. Incumbents on the other hand, often lack agility, cultural fit, and understanding for FinTech companies, are held back by internal constraints, and struggle to find internal executive sponsors because of their poor ability to articulate value proposition. A successful cooperation would have C-level sponsorship from the incumbent, shared objectives, and a vision. Incumbents would leverage on FinTech technology where they are blocked by their legacy systems, and FinTech companies would focus on a niche where incumbents underserve or are not profitable. There should be a mutual assessment of finances, technologies, business plans, and involved people. Thereby, incumbents should focus further on creating a culture that is more aligned with FinTech organisations, become more digitally agile, and focus their operational excellence on maximising customer focus. Also, they should build up an ecosystem for partnerships and strive for leveraging on big data insights, innovating with customer experience and satisfaction in mind, and focusing on empathy. Through this partnership, they should change their focus on hiring more relevant talents, and creating a culture that engages them.

While the authors agree with the direction of these conclusions, these should be extended through a strategic, an innovation, and an organisational perspective. Strategically, incumbents should first act as a FinTech company in the ecosystem of TechFin organisations, to understand better the direction of user expectations and limits of technological capabilities. Second, they should create their own ecosystems, leveraging on FinTech companies and other industry partnerships like co-branded credit cards, or loyalty cards. Third, they should use their position of power to block TechFin companies

wherever possible, and lobby for open standards. TechFin companies strive towards locking-in their users, and making it hard to leave the respective ecosystems. Unlike Australian banks that did the wrong thing with blocking innovation by taking a stance against Apple Pay, incumbents should rather drive innovation, however with the focus on connecting ecosystems and offering alternative, even better user experiences. The lock-in is the weak point of all TechFin companies, and should be exploited on all costs. Also, data protection and privacy should be the main strategic directives from incumbents, since public debate exposes TechFin, especially GAFA at the moment. As apps were key for smart phones to take off, FinTech companies could be key for alternative finance ecosystems. To increase innovation, incumbents should bring FinTech companies into their own ecosystems, invest in them, create their own FinTech start-ups, and work with them on alternative ecosystems without a TechFin lock-in. By providing FinTech companies with security and stability, this could move more of them into the incumbents' ecosystems, since TechFin companies tend to copy innovation from smaller Tech companies, which often leads to their insolvency. For example, Apple is known for embracing features that innovative companies created, and some apps they even banned just to copy their features shortly after (Kateliev 2017). Providing FinTech companies with a safe haven and an ecosystem could be beneficial for them and for incumbents. However, to fully leverage on strategy and innovation through interacting with FinTech companies, incumbents should analyse new start-ups, build clusters of their products, and use psychological knowledge to extract the underlying rationale of their functionality, which should act as a roadmap for future disruptive innovation and strategic positioning.

Additionally, disruptive innovation can also take place on an organisational level (Wan et al. 2015). Therefore, incumbents should deploy People Analytics, psychometrics and research on organisational culture for making data-driven decisions about aligning their organisations with their strategy and innovation. For example, Schneider (1987) found that employees are attracted (A) towards an organisation because of perceived resemblance of personality. If this perception is mutual, they are selected (S), however, if this mutual perception will not hold, there will be attrition (A). This ASA model explains why different companies have different cultures, it uses psychometrics to measure personality, and treats the individual employee as the smallest unit of corporate culture. Using this logic, Schneider and Bartram (2017) explored corporate performance and aggregate personality measures. They found that the personality trait Conscientiousness alone explains all variance.

However, they found in some iterations of their calculations that agreeableness could have an influence, as well. This may be aligned with different competency (Bartram 2005) demands at different positions, however research on aggregate personality is sparse. Future research should focus on aggregate personality on different job roles, and expand this line of research on business cooperations or joint ventures from the financial industry, as well as on future strategic positions where certain personality attributes may contribute to performance. For example, higher degrees of the personality trait openness in the top management team may re-align the strategic focus on innovation and organisational change (Judge et al. 1999; Oreg 2003).

While such deployment of psychometrics is rather found in recruitment and learning & development functions, the authors see it in People Analytics functions in the future, since this is where all quantified HR measures should be collected. Most corporations established such People Analytics departments already, comprised of psychometricians and computer scientists. These use all available and legally deployable internal and external data to create predictive models and potential scenarios of employee behaviour. Potential use cases are cultural fit, identification of ideal personalities for specific teams or tasks, prediction of attrition, performance, optimal career paths, group dynamics, risks, or recruitment. Only through this connection of internal and external data, incumbents can create a holistic understanding about their current and future needs both on an organisational level, yet also for better expanding into the life-spaces of their customers to better server their needs. In the specific case of psychometric tests however, it is standard to only get aggregate scores and final test results from providers or market research companies, which takes away a lot of granularity of the predictor space and therewith potential insights for organisations since it diminishes predictive capabilities. Therefore, the authors conclude that incumbents and TechFin companies alike should demand the raw data from their psychometrics providers. Also, they should cooperate with their People Analytics departments to better understand how to expand into various life-spaces of their customers, since these departments already gathered extensive experiences and professionalism in combining data science with psychometrics for optimising human behaviour. While most TechFin companies already established People Analytics functions, not all incumbents did (forthcoming working paper from the authors). Instead of focussing on products or product partnerships alone to create disruptive innovation, incumbent corporations and their FinTech partners could benefit from considering strategically, which talent they need to hire or develop for times that demand more flexibility, and especially for interface functions between the organisations. Also, HR

departments of incumbents need to be radically transformed and become as quantitative as marketing or finance. The authors predict the rise of People Analytics as the dominant force in HR departments within the next five years, and the radical change of the HR profession within the industry. Future HR departments will compete with other departments for the best quantitative talent, and will no longer be comprised of persons from subjects like the humanities. This is a relevant opportunity for specialised companies that either provide People Analytics consultancy, services, or applications for the finance industry. In its core, People Analytics supports the achievement of transformational goals that are aligned with overall corporate strategy. Such effects can also be achieved through internal innovation teams that act as pathfinders for the rest of the organisation. To align with the agility from start-ups, the authors further recommend incumbents to either invest and participate in start-ups or to establish internal innovation incubators that are independent from the rest of the organisation and have an executive sponsor on C-level. The goal is to enable these incubators to come up with innovation that is powerful enough to transform the bigger entity. Future research should focus on the deployment of People Analytics for improving innovation in the financial industry, the role of psychometrics in it, and how customers could become part of the extended development team, as TechFin companies do it with their current beta-test programs.

With regard to strategic innovation, the penetration of as many areas of the life-spaces of the customers as possible should be a goal for being able to compete with TechFin companies. Since incumbents are no technology provider, this should be achieved through strategic partnerships. Customers expect predictable behaviour that is the same across all products of the ecosystem, simplicity, and convenience. Also, products should improve social standing, or make them part of an attractive in-group. The underlying mechanism is bio-psycho-social. On a biological level, humans prefer low energy options, known patterns, and low stress situations. However, while they strive for a state of zero energy usage, homeostasis, outer factors urge them to adapt. If the adaptation happens in such a way that the behavioural patterns are only marginally changed, this change is accepted as positively. However deep-going change that could endanger adaptation, is perceived negatively (Selye 1955). For example, if the appearance of a banking app changes radically, includes new features, and promotes new FinTech partnerships, this could set off many clients. It would be much better to communicate all changes, include one feature at a time, and try to optimise the current user experience towards convenience, instead of totally changing it, unless it is the main pain point for customers. If a new experience is introduced, it should be explained

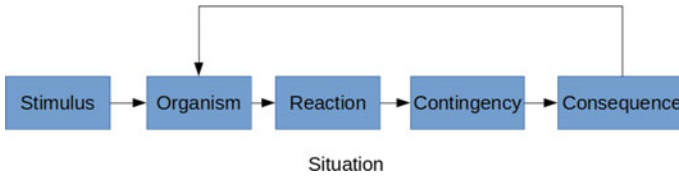


Fig. 1 SORCC Scheme of operant conditioning

thoroughly, optimally with a guided tour and preview versions. On a psychological level, habits are behaviours that are gradually built up over time in a specific place based on a fixed sets of clues and outcomes, and are always part of a bigger behavioural scheme. Through successful interactions, customers are conditioned to display the same behaviour, and internalise it over time. This can be explained with the SORCC scheme (Schneider and Weber-Papen 2010) (Fig. 1).

For example:

- *Situation/ Stimulus*—client is in the mall, smells fresh bread from a bakery, and uses the banking app to look for an ATM to withdraw money for shopping at a bakery. Yet the app is sluggish and does not deliver the utility needed.
- *Organism*—client is hungry and angry.
- *Reaction*—client clicks it away since it takes too long to load and opens Google Maps, which gives the answer in seconds, or even pro-actively, based on analysis of usual habits.
- *Contingency*—that happens in 70% of all times when client wants to use the banking app for something that other apps from TechFin companies can do, as well.
- *Consequence*—client learns that banking app is not convenient and stops using it.

Since this behaviour does only occur in 70% of all times, it is reinforced intermittently the effect is even stronger than if it would happen all of the time. Even worse, this prevents the customer from even giving future iterations of the app a try. In nature, behaviour does not result every time in an outcome, wherefore organisms display the strongest learning effect once it was successful, and then try to repeat it. This is one of the reasons why gamification works—however, its strongest effect is on a social level. Humans want to be part of an attractive in-group, and once in, they compare themselves with other members of that group. This effect is so strong that when

the wage of just one team member is cut, their performance decreases significantly higher than if all team members' wages were cut (Cohn et al. 2011). Also, persons tend to cheat more when they know that other team members earn more (John et al. 2014). Another social comparison process, which is the foundation of competition, is called mimetic desire - person A sees that person B wants something, and suddenly wants the same (Girard and Williams 1996). The green versus blue bubbles from Apple's iMessage system, which also allows the transfer of money from one user to another, manipulates its users with the same mechanisms: desirable in-group, definition of an out-group, and signalling of relative wealth (Masnick 2015). Over time, Apple iMessage and Apple Pay users are conditioned to prefer the payment mechanisms within the ecosystem that is perceived as not only more convenient and easier to use, but also as more prestigious, and it ensures continued membership in a desirable in-group. This is a direct threat to incumbents and FinTech companies alike. When humans are in a group, they influence each other, and create social norms through perceiving the behaviours and reactions of other group members. Thereby, distinct nudges help group members to identify the right behaviour, and the more social these nudges are the higher their effect on behaviour (van der Linden 2018). For example, an incumbent could partner with a FinTech company on an exclusive feature that signals high status, is only given to a select group, increases convenience, and offers clear in-group/ out-group signalling. Apple Pay was such a feature, however from a TechFin company. Though it helped participating banks, it put the other banks under pressure. For example, Barclay's did not offer Apple Pay in the UK in favour of their own solutions. They revised that decision after nine months since criticism from customers grew (Kelion 2016; O'Malley 2015).

With regard to disruptive innovation through FinTech, ideas and technology trends from FinTech companies could be incorporated into new product development of incumbents, or they could cooperate with FinTech companies with the right organisational fit. The authors identify the following FinTech start-ups as exemplary for current FinTech trends.

- *Trasferwise* from Estonia reduces cross-country transfers by keeping the cash flow inside each country through matching people. They received funding from Sir Richard Branson. The basic concept is that of a social network through peer-to-peer money transfer.
- *Revolut* is a London-based company that uses social networks to exchange money internationally through peer-to-peer transfer, and offers pre-paid

debit cards. Customers do not pay for most of the services, but a mark-up on interbank exchange rates applies.

- *Monzo* is a London-based banking company that only exists on the smart phones of customers that build together a community of users. They offer financial planners and notification services around the needs of customers.
- *Faircent* is a peer-to-peer lending service from India that uses machine learning to match lenders and borrowers
- *Kissh* from India provides instant credits for customers at digital purchases.
- *Chasing Returns* from Ireland helps traders to overcome their psychological biases.
- *Trezeo* from Ireland helps entrepreneurs to stabilise their income and make savings.

These alternative finance companies are mobile first, convenient, and follow the design and user experience rules from the platforms they're operating on—Apple and Google—which enables a consistent experience. Their products are based on the psychological effects of social networks and social norms, and they use machine learning to operationalise new sources of data. This enables them to provide niche solutions that incumbents cannot or do not want to provide. Also, they operate on low budget and therefore have to focus on frugal innovation (Prabhu 2017). In summary, they innovate disruptively by creating new products or services that replace a niche of the existing offering, are cheaper, are good enough to attract more customers, and finally force the incumbents to react (Wan et al. 2015) or to cooperate. Beyond mobile first, the authors deduce three main trends with FinTech companies:

1. Increased personalisation and convenience through data insights
2. Behavioural modification through nudges or visualisation
3. Social processes like group-lending or crowd-funding

First, with regard to personalisation, social log-ins from platforms like Facebook offer increased convenience. When users log into services using social log-ins, data is exchanged between the social network and the service. This could enable new better applications for alternative finance companies. For example, a social log-in through Facebook may augment or replace the decision-making process at a loan application. Data from social log-ins may provide information like the financial stability of the network of friends, risk level of hobbies, predicted life success based on intelligence (Wei and Stillwell 2017), life stability risks (Park et al. 2015), and personality (Kosinski

et al. 2015). The personality factor Conscientiousness is one of the strongest predictors for professional success, Neuroticism for stability in difficult times, and Agreeableness for getting along with others (Schmitt 2014). This data could be combined with additional sources. LinkedIn may provide proof that a person has a stable career and receives positive feedback from others. Location and health data from smart phones and smart watches could provide information about a healthy and safe lifestyle. Sharing data with alternative finance companies could be facilitated through convenience benefits and data insights for customers like easier opening of accounts, better personalised and maybe even predictive service, more bespoke offerings, monthly trends, general spending patterns, and lifestyle advices. The trend to exchange data for convenience, price reduction, or additional services is already established and accepted by consumers in various industries.

Second, with regard to behavioural modification, data sharing comes first to understand the baseline of an individual user, which is then used to personalise the services even more. This, in turn, may nudge customers further to share even more data. For example, insurance providers like Vitality offer insurance fees reduction for customers who share data from their fitness trackers (Cameron et al. 2018). Loyalty apps like that from Starbucks offer convenience features like pre-order, which nudges people to use these apps, share data, and enjoy the benefits it brings them. The hospitality service Airbnb uses offline and online verification. They demand its users to scan their ID cards and to connect their Facebook and LinkedIn with their Airbnb accounts. This approach reduces administrative effort and increases the level of mutual trust and convenience (Geron 2013). Dutch carrier KLM offers the option to chose a seat neighbour on a flight based on each others' Facebook or LinkedIn profile (Clark 2012). While fully automated decision-making about persons is forbidden through the General Data Protection Regulation (GDPR) in the European Union (European Commission 2018), European alternative finance companies could use new sources of data to augment human decision-making, for increasing the speed and convenience as well as the relevance for their customers, which, in turn, would nudge them to share even more data and use even more offers. Some banks in the US already started using social media to inform their decisions on loans and pricing (Quittner 2012).

Third, with regard to social processes, gamification and social comparison with friends to compete for an overall better score and therewith membership in a prestigious in-group could lead to more data sharing. Social processes are already used by a variety of industries to encourage increased interaction as well as more data sharing. With regard to group lending, sufficient

data about individuals and how their behaviours can be changed must be collected. The authors believe that social media-centric offers yield opportunities for opening new market segments for financial service companies. Offers for customers could be created, who may have a low credit score, but are embedded in a strong network that could hold them responsible, warrants for them, and could support them financially. The principle of lending to a group and not to an individual is fundamental to micro-credits for the poor, which have been pioneered by Nobel Price winner Muhammad Yunus. Applications for alternative finance could be based on peer-to-peer lending and crowd-funding. A mobile-first and block chain-based approach could decentralise this offer, and match-making as well as decision support could be informed through social media and ambient data from other IoT devices. Future research could import principles of frugal innovation (Prabhu 2017) for alternative finance. For example, mobile payments were pioneered in Africa and China because of the lack of an existing infrastructure.

Also, future research on trust in data sharing should be conducted. It should cover the change of norms to better predict demand based on societal trends, and understand how and under which ethical imperatives organisations can influence such developments. Alternative finance research could benefit from importing concepts from marketing strategy, social psychology, and geographic psychology. Exemplary topics could be how to create personality profiles based on social media data and target people for mass persuasion towards a new trend like cashless payments (Matz et al. 2017), building up social influence to initiate a viral participation in crowd-funding for a good cause (van der Linden 2017), or predicting investment risks based on geographic distribution of personality factors that may lead to populist rises (Obschonka et al. 2018). The authors agree with the conclusion of the industry report from World Economic Forum in cooperation with Deloitte on the ‘The Future of Financial Services—How disruptive innovations are reshaping the way financial services are structured, provisioned and consumed’ (McWaters 2015):

- Innovation will take place where the most customers are and the most profits can be made
- The highest impact can be achieved through a combination of low investment, strong data collection, and focus on providing platforms
- While the banking sector offers low hanging fruits, the insurance sector yields the biggest potential for disruption
- Incumbents and alternative finance providers will compete for the same spot

- Incumbents, new entrants, and regulators will have to cooperate closely
- Disruption takes place over time.

In summary, incumbents could benefit from adapting trends from other countries or inventions of disruptive FinTech start-ups, and from cooperating with these. They should increase their efforts to accelerate the transformation of their organisations through adaptation of new technologies, deployment of People Analytics, and creation of innovation incubators. Future research should also focus on the top management team and CEO of these organisations, and explore the influence of their personality traits on innovation. Some progress has been made already. For example, Deutsche Bank now offers a service called 'Finanzplaner' (German, 'financial planner'), which allows their customers to integrate real-time data from other financial institutions in an overview of their financial situation. Investment banks started using social media data to predict the behaviour of people and how this influences stock prices. Cars in front of stores are counted through satellite images and allow insights in real business of potential investments. Also, the geographic terrain measured by satellite images is used to determine home values (Glaeser et al. 2008) in combination with public data from cities or Airbnb. This allows much more precise predictions of real estate prices. However, they should develop a sense of urgency, as well. Nokia, Siemens and Blackberry were challenged by newcomers, but it only needed one innovation from another industry, Apple's iPhone, to fundamentally change their business model. The U.S. Department of Justice assessed Microsoft for potential counter-competitive actions, and they found internal Microsoft documents that described their strategy of 'embrace, extend, and extinguish' (Economist 2000). First, they embrace a new market like open standards for office suites, then they extend its products in a way that only Microsoft products can be used, just to finally extinguish competitors through this. In variations, this approach has since been copied by various Silicon Valley companies. That's one of the reasons why European car makers acquired NOKIA's map services. The future of automotive will be self-driving, the car makers don't want to experience what happened with the mobile phone industry. Amazon is the company that is imposing the biggest threat to incumbents, and incumbents should have an answer to that, if they don't want to face their 'iPhone' moment. Psychology and psychometrics together with machine learning could open new and alternative strategies that are beneficial for incumbents and for FinTech companies alike.

2 Deployment of Machine Learning in Alternative Finance

Current technological trends and developments are based on psychological effects that are best to be described by constructs like social norms, in-group preference, trust, personality, ability, motivation, performance potential, competencies, propensity towards certain actions, values, and happiness. However, there are no surveys, measurements or questionnaires in place, but most existing offers either trigger psychological needs like the urge to help others, predict psychological latent traits by machine learning models, or they use psychological factors and theories to engineer features for machine learning models. This fusion of psychometrics and computer sciences has only just begun, and it defines the core of alternative finance FinTech start-ups. Though not obvious, it is important to understand that the principles and rationale behind these start-ups and the incumbents are fundamentally different. That's nowhere as observable as in the methodology deployed that powers the products. While the foundation of the established financial industry still is based on microeconomics and econometrics, the challengers deploy machine learning, deep learning, and artificial intelligence to pursue their goals. They are at the forefront of the fourth industrial revolution, in which technology and humanity grows even further together. This fusion is foremost mirrored in the change of the methodologies deployed. While still valid, old econometrics models will cease to suffice, and will need to be augmented with machine learning models (Varian 2014; Mullainathan and Spiess 2017; Athey 2018). This marks the point when also theory or industrial knowledge ceases to suffice, and data-first methodologies will need to be deployed. This, in turn, demands a strong sense of ethics and responsibility. The authors extract the following three topics of this evolution of methodologies:

- Restrictions of economics
- Growing importance of deep learning
- Responsibility

2.1 Restrictions of Economics

The prevailing logic of economics and econometrics is that of industrial knowledge, which helps forming a theory about the data generating process behind collected data. Based on a strong theory, and *ceteris paribus*, individual effects are identified and a model of this process is created. These models

are of probabilistic character, which means that quality measures and significance levels are described that indicate how good these models describe and predict observed reality. Thereby, lack of perfect fit is attributed to measurement errors, omitted variables, lack of data, hidden connections between variables, or lack of industrial knowledge. This usually comprises the error term. Also, a number of assumptions and tests whether these are met need to be done, and the models have to be adapted based on these outcomes. In short, economics and econometrics follow mostly a theory-driven top-down approach but of certain exceptions like ‘atheoretical’ (Box and Jenkins 1976) Arima time-series models. However, also these work by—mostly manually—‘matching behaviors of the sample autocorrelation function (ACF) and partial autocorrelation function (PACF) to the theoretical autocorrelation functions’ (Ngo 2013, p. 1).

Let’s consider the above for a moment. Data is produced on an unprecedented scale, and most of it is based on human behaviours and interaction with technology. Although some of these data points may be connected with direct theory, others are part of a bigger pattern, in which existing theories may or may not unfold. For example, if a person applying for a loan lives in one part of a town that is not considered to be wealthy, but spends most of the time in a wealthy area according to the geolocation data from the smart phone, that might have different reasons (see BJORKEGREN and GRISSEN (2018) about a real study on smart phone data and loan payment). Either the person lies with the address, or the person works in the wealthy area. Now, this person may spend rather nights than days in that area, which may indicate a night job like security, or on the other hand it might be indicative of a love affair in that area, which would of course increase the financial power in that person’s network. Let’s further assume this person starts posting more about holidays and chats about buying new electronics since the nightly habits changed. This could be indicative of being a thief or further deepen the assumption of a love affair. Maybe at the same time, the number of wealthy friends and followers in social networks increased, and the banking app registers a new luxury phone. This could be further indicative of a wealthy relationship, and thus over time increase the odds of being rated with a higher credit rating. As an interim step, one may think of testing these assumptions with interjecting advertisement for luxury goods and romantic presents in that person’s network traffic. If a person buys some of these, further indication of a change in relationship is given.

This may sound like a trivial and obvious theory, but there could be countless different explanations and unknown factors that become obvious only once one either observe a large enough number of different cases, or when

more about the context of an individual is known. Also, there might be some factors that are only understandable with a thorough command of psychological theory. For example, a hospital may wish to increase patient experience because that allows them to demand higher prices or simply improve treatment costs due to psychological factors. The hospital may use the sentiments in the online postings and emails of their patients to better understand the polarity of their language and thus estimate the patient experience. However, clinical psychology knows the concept of secondary gain. So, from a certain level of experience onwards, that person might want to stay longer in the hospital because of the good service and additional attention from relatives, which would, in turn, increase the treatment costs. However, the extent of this effect is not known, as is not its expression through gender, age, and ethnicity. Also, it is not known whether culture, personal values, or personality may have an influence on that. One may assume that conscientious persons would not stay longer, however this could change depending on their happiness in marriage and the dominance of their significant others on their life decisions. This example alone shows that theory is limited by the number of options and unknown possible root causes, especially with regard to human action and interaction and their respective environments.

2.2 Growing Importance of Deep Learning

Therefore, a much better approach would be to create new synthetic variables based on knowledge about psychology and human interaction, add them to existing variables, and use computer algorithms to find such patterns that occur often together and predict desired outcomes like creditworthiness or loan default risks. This is exactly what machine learning algorithms do. In older machine learning approaches, which are still inspired by statistics, the creation of new synthetic variables is often called ‘feature engineering’. However, in modern deep learning approaches, which are more powerful and rather inspired by cognitive neuroscience and psychometrics, feature engineering is the process of deriving abstraction from the input that is best linked with the output. This approach is faster and computationally more efficient than classical machine learning approaches.

Machine learning can roughly be classified in supervised, unsupervised, and reinforcement learning. Supervised learning is based on experts telling the algorithms, which feature should be connected with which output. For example, most neural networks belong in this category, but also generalised logistic regressions, if the feature selection process is mostly automated. The reader should think of it as feeding in as many features as possible, and

realising the final model through the outcomes of the selection process. Artificial neural networks are heavily influenced by findings from neuroscience and psychometrics. They simulate the way the brain processes information through various computational layers. Thereby, information enters the artificial neural network through the first layer, called input layer, gets computed through the internal layers, and results are given out by the last layer, called output layer. Typical applications would be pattern recognition, for example when reading remittance orders filled out by hand. Supervised learning is conceptually most close to the classical econometrics approach, and authors with an economics or econometrics background give stronger emphasis on these algorithms, when they introduce machine learning (Varian 2014; Mullainathan and Spiess 2017; Athey 2018).

However, rather empirically oriented authors focus more strongly on unsupervised learning, for example Gu et al. (2018), who use it to measure asset risk premia. Unsupervised learning is basically an algorithm that looks for patterns in the data and learns by itself, which combination of features is best to predict which outcomes, respectively, which outcomes are best suited for which set of features. Thereby, the algorithm does not know the outcomes upfront, but it derives features that compress the data in the best possible way. These features are then fed into supervised machine learning algorithms. This is for example used in sentiment analysis with neural auto-encoders, which feed their output in shallow neural networks like word2vec models that create word embeddings. Another example is some forms of Random Forest™ models that deploy synthetic data based on a reference distribution to create dissimilarity with unlabelled observed data (Breiman 2001). This could be the communication data of persons with high credit ratings, which could be used as a blueprint for simulating a dataset for a Random Forest™ algorithm and compare it with communication data from loan-seekers.

Reinforcement learning focuses on finding outcomes for problems where experts cannot supply feature-outcome pairs upfront, but where the algorithm acts like an agent. It interacts with an environment, discovers own examples, and learns through experience. It has to discover how the environment works, and obtains data by itself as it explores. This is inspired by human and animal behaviour, which takes place in specific environments, is goal-oriented, inhibited by punishment, and facilitated by rewards like Pavlov's dog. Through conditioned experiences, future behaviour is shaped. For example, if a child accidentally touches an oven and removes the hand in pain, it would learn to stop this behaviour for life. In reinforcement learning, this would be called a training pair. A potential application could be a self-learning dialogue system for call centres in banks. They interact

with customers, gather experiences, and understand successful behaviour through reward signals like quick solutions of customer satisfaction. Google's AlphaGo, the first computer program that ever beat a professional human Go player, deployed reinforcement learning, as well (Silver et al. 2017). A comparable algorithm could be developed to do autonomous investment decisions in complex environments and learn over time to beat human investors. However, as with self-driving cars, trial-and-error might be risky in complex trading situations. Both, human-machine interaction systems and reinforcement trading algorithms yield strong future potential for alternative finance, and the authors recommend further exploration.

Deep learning is the dominant approach in the machine learning community due to its computational efficiency, robustness, and ease of application. This trend arrived at the finance industry, as well. Hedge funds like Citadel (second biggest hedge fund in the US) started deploying deep learning for algorithmic trading and identifying long term investments. They hired a former director of AI research labs at Microsoft to help them realising this change. One of their most important applications is natural language processing based on data from twitter and other social media. This is insofar relevant, as hedge funds were previously in favour of established statistical models due to the legal risk involved. They have to estimate the risk of their investments for insurance purposes upfront, and most deep learning models are more like black boxes, and there is only little theory to predict their behaviour. Therefore, the estimation of risk is problematic. However, the success of deep learning is so much better than theory-driven approaches, that they decided to deploy it wherever possible. Since investment banks have to be on top of the curve to be competitive, they started deploying neural networks in their trading software for internal uses, as well. The authors strongly recommend deeper studies of deep learning for alternative finance practitioners and suggest future research on how deep learning can replace theory-driven econometric and machine learning approaches. At the same time, the authors acknowledge the importance of ethical boundaries, and personal accountability of data scientists. Therefore, they stress the additional importance of further research on methodologies that allow observing the inner mechanisms of deep learning models to avoid biased decision making and currently unforeseeable financial risks alike.

2.3 Responsibility

Google's selfish ledger (Savov 2018), could be realised with deep reinforcement learning, and applied ethically, assist its users to reach their goals. There

is already a trend for artificial behaviour modification that is embraced by customers. For example, FinTech start-ups offer help their customers to optimise their individual finances based on behavioural changes. MindBlown Labs uses games to teach Millennials retirement plans. Based on gamification, behavioural theory, and learning from experienced wins and losses, desired new behaviours are reinforced. This shows how much machine learning and psychometrics are intertwined, and that the nascent fourth industrial revolution will have strong influence on modification of human behaviour. Actually, one of the main, if not the main figure of modern deep learning, Geoffrey Hinton, was first a psychologist, who then turned into a computer scientist. The authors predict that new subjects will emerge that connect psychometrics and artificial intelligence research, and that these subjects will become part of the basic education of finance experts. Alternative finance research should be directed on that focus, and potential FinTech start-ups should explore emerging opportunities.

However, this development with its opportunities bears significant reputational and legal risks. The most advanced algorithms almost unanimously operate like black-boxes. That means that though they provide superior predictive power over established econometrics measures, many of them do not produce a visible model but only input–output relations. While this may go well in most of the cases, pure reliance on algorithms without human intervention or theoretical input may lead to false classifications or predictions that may yield undesired consequences for organisations that deploy them. This is especially true, if data integrity and quality is neglected in favour of putting all available resources into creating deep learning algorithms. For example, Microsoft created a chat-bot that was turned over by internet trolls into a racist in short time (Vincent 2016). Or, Amazon stopped a ‘secret’ internal recruitment system that preferred men (Dastin 2018). Based on the sheer number of male candidates in the technology industry in historic data, the probability of a man being rated as a top candidate in the past was much higher, wherefore the algorithm was trained on pre-biased data. This bias is one of many reasons why the pure reliance on algorithmic decision-making is forbidden in Europe due to the GDPR (European Commission 2018). Other reasons at the root cause of regulations are algorithmic challenges that have strong ethical implications. For example, since facial-recognition technology still struggles with correctly identifying people of colour, Google blocked a whole group of picture recognition algorithms, that still classify often persons of colour ‘gorillas’ (!) (Vincent 2018). Alternative finance technology that relies heavily on facial recognition like Apple Pay or on video analysis like the Amazon stores, should carefully balance the deployment

of available technology with reputational dangers and business strategy. In case of doubt, conservative strategies may yield the overall higher benefits for growing one's business and defending the competitive edge. This challenge goes well beyond the proverbial 'garbage-in, garbage-out' analogy that many data scientists use. Prevalent social stereotypes and unconscious biases are deeply rooted in human nature, therefore a strong ethical stance has to be developed, which, in turn, will help building up trust with customers and further benefit growing business. The authors conclude that alternative finance FinTech start-ups would benefit from basing all of their decisions on strong ethical guidelines. Future research should explore in-depth algorithmic ethics for alternative finance, especially on the overlap of artificial intelligence, psychometrics, and business strategy.

3 Ethical Risks, Regulations, and Opportunities

Beyond engineering problems and biased data, the rationale behind algorithms may fundamentally differ depending on the background of their architects. For example, self-driving cars would have to decide in worst-case scenarios whom to kill if they had only two choices. A recent influential paper in *Nature* explores this dilemma and finds significant cultural differences on the question whom to spare. Southern cultures spare young persons, females, and higher status persons, Eastern cultures spare lawful citizens and pedestrians, and Western cultures spare larger groups, though generally prefer inaction, which means taking no choice at all (Awad et al. 2018). This yields implications on the behaviour of such cars if they are developed by companies from different cultures. Would a Hyundai decide to kill a robber? Where would the car know it from? Would the data be freely available, which would be equivalent of a kill list? Would a BMW rather kill a Nobel Price winner instead of a group of rapists? Again, where would the data come from, and shall corporations have such data at their disposal? Also, car manufacturers may be tempted to offer different behaviours in different cultures to protect their reputation and the trust in their brand. This may be perceived as hypocrite and could damage their reputation if it was ever leaked. Even worse, culture-specific algorithmic behaviour challenges the fundamental universal applicability of ethics. It breaks with established ethics based on Asimov's first law of robotics that a robot shall never harm human beings. Completing this thought would result in fully autonomous killing machines, an outlook, most artificial intelligence (AI) researchers strongly oppose (Walsh 2015). A

financial AI that influences stock markets may impose even stronger existential threats than autonomous killing vehicles. For the sake of profit, a machine could outsmart the markets, level up food prices, destroy jobs, launch disinformation campaigns to initiate wars, buy too much energy until the power grid of a country breaks down, or crush the currency of a whole nation—all of that to maximise revenue that it doesn't even understand or can make use of. Although such future scenarios are important, valid, and discussed (Bostrom and Yudkowsky 2014), the authors identify six issues that affect alternative finance today, and which should be addressed in the current debate in alternative finance:

- Trust
- Acceptance
- Privacy
- Freedom of choice
- Reciprocity
- Transparency

3.1 Trust

Ribeiro et al. (2016) develop a new technology that helps interpreting the results from various models like Random Forest™ for text Analysis and artificial neural networks for image classification. In one of their experiments, they try to understand why huskies are wrongly classified as wolves. In principle, their approach is to identify and display each element that leads to a sub-decision, and visualise it. The result in the husky-wolf experiment is that the presence of snow is the strongest differentiator. This skews the classification decision of the algorithm towards wolves strongly. Research on understanding 'black-box' models is in early stages and has yet to arrive in most business applications. Ribeiro et al. (2016) argued that 'trust is crucial for effective human interaction with machine learning systems, and that explaining individual predictions is important in assessing trust' (p. 10) (Fig. 2).

Human beings need trust to be fully productive and functional. Once trust is broken, companies not only lose a client, but a network of clients. For example, the most important task of Apple Store Retail employees are to empathise and repair trust, and not to sell products (Bunning 2018). And, 'just' a bad recruitment experience may damage not only the bottom line, but also the consumer brand (Baker 2010), which may summarise to hidden losses in the millions. After the Cambridge Analytica scandal, user trust in Facebook declined by 66%, which led to a decline of 199 billion USD market



Fig. 2 Dog classified as wolf

cap and a plunge in share of 19% (Neate 2018; Weisbaum 2018). While the Obama campaign team used Facebook data collected from their app, as well, the main differentiator is that people were asked and voluntarily give data, whereas the data from Cambridge Analytica was generated under false pretence. One action was mutually agreed upon, the other was a major breach of trust (Fig. 3).



Fig. 3 Explanation

3.2 Acceptance

The next problem occurs when machines start making decisions in a way that makes humans feel that their freedom is restricted. This feeling could be based on wrong communication or a decision-making process that cannot be influenced. The natural human reaction to this is to display behaviour that shall re-establish the perceived lost freedom through indirect means. This psychological phenomenon is called reactance (Brehm 1966). For example, if the credit limit of a bank customer is changed, and if this change is not explained or announced, the customer may open a credit card account with another bank. Since algorithms may make such decisions without further announcement, this could be a potential hidden source of declining revenue. Without understanding psychological theory and psychometrics, these root causes may stay hidden, and measures would only treat symptoms. Another problem is the so-called uncanny valley from social research on human–robot interaction (MacDorman and Ishiguro 2006). A robot that is too close to human appearance, may evoke strong feelings of repulsion and fear in humans. This is the reason why most robotics researchers focus on the cuteness of robots while maintaining a level of anthropomorphism (Rosenthal-von der Pütten and Krämer 2014). The authors propose to expand this phenomenon towards virtual assistants and any other kind of computer–human interaction, for example in automated phone systems for telephone banking. While existing systems that offer pre-defined options for the caller are established, it would be important to understand what effects systems have that perfectly simulate human communication. For example, Google’s Duplex technology is an artificial intelligence, which makes real calls for various standard tasks like reserving tables at restaurants. These calls are so realistic, that they cannot be identified as a machine. The developers included breaks, random sounds, differences in tonation, and emotion in the voice of it (Welch 2018). While this advancement is remarkable, the reactions on it are negative—especially since callers would not know that they are talking with a machine. Therefore, Google plans to let the public available versions identify itself as such. Developers of alternative finance products that simulate human interaction should know these human reactions, for example to chat-bots or virtual assistants, and implement measures that prevent adverse reactions that yield potential business or reputational risks.

3.3 Privacy

Another consideration should be taken about which cooperation partners to chose and where. For example, as soon as it emerged that Facebook tried to access data from banks, which included account balances and spending habits, a public outcry resulted (Harwell 2018). Though they genuinely want to enable users of their messenger service to check their accounts via chat and thus offer banks new customers who value this service, this effort backfires due to their bad reputation with regard to data protection and ethical behaviour. The public remembered a series of severe breaches of trust. In 2010, they conducted without consent an experiment by displaying voting advertisement to circa 61 Million users to increase voting behaviour (Bond et al. 2012b). In 2012, they conducted without consent an experiment on inducing emotional contagion. For this, they altered the emotions of users by manipulating their news feeds. While being published in the Proceedings of the National Academy of Sciences of the United States of America, the editors express their concern and correction upfront (Kramer et al. 2014). In 2017, they experimented in Sri Lanka, Bolivia, Slovakia, Serbia, Guatemala, and Cambodia with changing their news-feed and erasing official news to varying degrees. In Cambodia, this had an influence on the unstable political landscape and undermined the efforts of democratic forces and propagated fake news (Paviour 2017). The list of privacy invasion, trust abuse, scandals, suspicions, and rumours is too long to report here, and it reaches from shadow profiling (Brandom 2018), abuse of phone numbers for two factor authentication (Lomas 2018), Cambridge Analytica (Cadwalladr and Graham-Harrison 2018), participation in the PRISM program (Nolan 2013), unhinged sharing of data in its early days (Roose 2018), sharing data with the Chinese companies that are close to the Chinese government (Shepardson 2018), funding through the CIA (Edwards 2011), to Mark Zuckerberg allegedly wanting to run for office (Carter 2017). In India, the Facebook company WhatsApp was the tool of choice to organise a lynch-mob (Satish 2018). More than that, the use of Facebook decreases psychological well-being (Shakya and Christakis 2017) and has a stifling effect on minorities, which undermines the democratic process (Stoycheff 2016). In summary, the behaviour of a company, whose CEO named his customers 'dumb fucks' for trusting them with their data (Orlowski 2010) is the blueprint of everything an alternative finance start-up should not do, and cooperation with Facebook and its affiliates should be as transparent as possible to prevent a spillover and loss of customers.

Facebook has no payment service yet, while Apple, Samsung, and Google do. This is despite the fact that Google openly analyses as much data as possible from their customers. They even experimented with hands-free payment that transfers the pictures of people standing close to the counter to the cashiers for identification (Heater 2017) and struck a deal with MasterCard and other providers of credit and debit cards that enables them to connect online advertisement with offline spending habits. They state that through such business partnerships, they have access to the data of 70% of all US Americans with such cards (Bergen and Surane 2018). Consumers trust them and not Facebook, because of data safety and open communication. However, the perception of Facebook depends on the region and the experiences customers made with them. For example, it helped the organisers of the Arab Spring in their struggle for liberation. An Egyptian father was so grateful that he named his daughter 'Facebook' (Abdelaziz 2011). In his case, the benefits far outweighed any possible privacy concern. People in Estonia seem to enjoy the benefit from their e-passports, as do people in Iceland, where most of the contracts, including gym membership, demand the social security number of a person, a situation similar to Finland. However, in the aftermath of the Cambridge Analytica scandal, up to 300,000 Finns erased their Facebook account due to privacy concerns, which is a huge number for a country of only 5.5 Million people (Teivainen 2018). This shows that a breach of trust can have a devastating effect on businesses, and that respecting the privacy preferences of customers is crucial. The authors see potential in further research on privacy and data protection, and how this can be used for ethical enrichment of financial products, customer retention, and word of mouth marketing.

3.4 Freedom of Choice

People will accept new technologies better if they are not forced to using them. The speed and convenience of credit cards is best marketed if the option of cash exists. Only companies with a strong pull-factor like Apple transcend this acceptance mechanism and convince their users with the strength of their brand and the scarcity of their product—based on limited availability, price, and being part of an 'elitary' in-group. Technically, Apple acts as a provider for alternative finance services though they are an established incumbent in the computer industry. Their cash reserves of 285.1 billion USD (Balakrishnan 2018) would place them on place 42 on the list of GDP per country—right between Pakistan's 304 billion USD and Chile's 277 billion USD GDP in 2017 (International Monetary Fund 2018). Their

brand leads the Forbes list since eight years, has a value of about 182.8 billion USD—about the GDP of Algeria (178.3 billion USD) and Madagascar (10.6 billion USD) united (Badenhausen 2018). The mere vision of their logo and products activates the same brain centres of some of their customers as that of religious people when seeing religious iconography. Furthermore, they display all elements of a religion - a clear and big vision, a sense of belonging, an enemy, a distinct narrative, a foundational myth, word of mouth, strong symbols, and rituals (Lindstrom 2008).

For example, their influence on consumers is so strong that Apple Pay is pervasive in Japanese society, even in the rural areas. In Tokyo, the majority of people pay with their iPhone or Apple Watch for daily needs like convenience stores, subways, buses, and taxis. Neither conservative banks in Australia could uphold their opposition against the desires of their customers to use Apple Pay, nor were home-baked alternative solutions from Barclays in the UK successful. While Apple Pay was just rolled-out with a selection of banks in cash dominated Germany, one major German bank, 'Sparkasse', created their own solution to offer an alternative since they perceived the offer from Apple as too one-sided (Mühlbauer 2019), and they would lose direct access to the data. The authors predict that this solution will be either replaced or strongly inhibited by Apple Pay, and that 'Sparkasse' will be forced to accept Apple Pay at one point in time. This assumption is based on the market dominance of Apple in the upper-income category, the brand loyalty of their customers, the market entry strategy based on scarcity that is comparable to that of the iPhone, the marketing message of privacy that resonates well with German preferences (Guenole et al. 2015), the example of Barclays in the UK, and the existing infrastructure for contactless payment. The existing infrastructure is key to the success of Apple Pay in Japan, as well, where even old vending machines support this system, since Japanese iPhones come with a special chip that allows them to use the established SUICA system. Other companies try to copy their style, for example Xiaomi. Its founder and CEO Lei Jun sometimes even copies the clothing style of Steve Jobs.

It has to be seen whether Apple deserves their customers' trust. In spite of denying it, there is indication that Apple took part in the PRISM program (Greenwald and MacAskill 2013) and cooperated with governments on data exchange (Smith 2017). iPhones may load meta-data to iCloud (Yoo 2016), which makes it much more accessible if it would have remained on the device. They moved iCloud data, including encryption keys, to China (Nellis and Cadell 2018), and banned VPN from the Chinese app store (Nichols 2017), which were used to circumvent censorship. Recently, they hired Google's former head of AI, John Giannandrea, to run their machine learning and

AI strategy group. The influence of this group is visible in the way products make more on-device use of their users' data (Nicas and Metz 2018). Until now, Apple is the only big technology incumbent that has privacy in the core of their strategic positioning against their competitors, and they rely on selling fewer products at a higher margin. A leak about abusive data collection would mean an distinct breach of trust and severe risk for business continuity. Therefore, the authors assume that Apple's efforts are mostly genuine and that there will be mechanisms of plausible deniability for adverse scenarios. Apple offers the choice to share data or not and is quite vocal about it. This is linked to their success, and copying Apple's focus on privacy could be used for ethical enrichment of products in alternative finance.

3.5 Reciprocity

A fruitful idea could be as well to copy Apple's early strategy, when they positioned themselves as the rebels against the established computer industry. Back then, they used ethics and out-of-the-box thinking to come up with innovations that disrupted the market: giving the personal computer to a broader group of consumers and thereby empowering them. Comparably, this happens today in the social media business with start-ups like minds, BitChute, LBRY or diaspora. While the incumbents Facebook, Twitter, and YouTube become more mainstream and cooperate tightly with media agencies, these young companies focus on freedom of speech and giving broadcasting to a broader group of consumers. However, they are under attack because they attract fringe groups and extremists. Alternative finance provider Patreon just suspended BitChute for extremist content. While their argument that BitChute provides a platform for extremists is true, so do Facebook and Youtube, before such content is uploaded and not filtered. While BitChute argues with freedom of speech, and Patreon argues with societal responsibility, Patreon may miss out the opportunity to become the go-to-financial service provider for alternative media, especially since BitChute and LBRY seem to be the only serious YouTube competitors. Another competitor, SubscribeStar, just emerges and tries to exploit this gap. The authors assume that a cooperative strategy with BitChute could lead to the desired effect of getting rid of the worst extremist content while still offering alternatives to the main stream media. Also, the better availability of hate speech filters like the Online Hate Index from ADL and University of California, Berkeley's D-Lab could become mandatory for such cooperations. This would allow alternative platforms to emerge, ensure their full cooperation with payment

providers, and keep the public safe from hateful and extremist content at the same time.

This reciprocity is not only important on a business to business, but also on a business to consumer level. Customers are willing to share data if they feel reciprocity—for example if they get a price reduction through their loyalty cards. In a study, significantly more people were willing to give out information if they received the reward for this before they handed over information—43.5 instead of 29.8%. The participants reported a significantly higher internal pressure to hand over information - 47.9 instead of 39.9% (Happ et al. 2016). The psychology of this is that people want to be socially accepted and part of an in-group, even if it is just as spontaneous as an interview. If someone invests in them, this is a sign of being part of the in-group, which triggers pressure to reciprocate. Another driver is striving for competition and achievement, which explains the success of gamification or point systems. There are indications that men respond better to competitive settings against other groups than women (VanVugt 2008). Since reciprocity is a strong social mechanism, it can be easily exploited to manipulate people. The underlying psychological drivers of in-group acceptance and competitive pressure are two primary human motivators: wanting or fearing something. Therefore, people rather comply and react with conformity if there is the risk of being rejected (Ash 1951). Game theory shows that people tend to avoid perceived risks (Kahneman and Tversky 1979), and therewith, rejection. Incumbents already trigger these powerful drivers in their customers to advertise their products. For example, insurance companies promise security against risks, and banks promise obtaining desired products through low-risk loans. This exchange is based on perceived rather than rational reciprocity. Alternative finance providers would benefit from knowing psychological research very well, and use a mixed strategy of conformity and non-conformity with communicated promises from their competitors, based on the market segment they want to address.

For example, people often trust hazard-related information more, if the message is negative. However, research indicates that this is strongly moderated by prior beliefs. If messages are more congruent with prior beliefs, positive messages have a higher impact on risk perception than negative ones (White et al. 2003). Group-lending FinTech companies may make use of this, and offer peer-to-peer insurance products. These yield the potential to disrupt the insurance market, however the narrative will be very important, since trust is not invoked by the name of a big brand, but by belonging to a group of people who secure each other.

Australian Huddle Insurance entered the market with such peer-to-peer offers, however their narrative focused on low costs. While they still use peer-to-peer technology, chat-bots, and artificial intelligence as customer interfaces, it is unclear what happened with their original peer-to-peer lending products. Future research could focus on that and on the change of their marketing message. Based on the rationale above, the authors assume that peer-to-peer insurance companies would benefit from focussing on a longer than usual marketing campaign before market entry. This should deliver a positive narrative about the beneficial effects of backing each other, rather than being insured by an anonymous company. The goal is to change prior beliefs, and only enter these markets when the probability of acceptance is high. This research could add substantially to the theory of cultural and economic factors of product take off (Tellis et al. 2003), and the authors identify strong potential for future research for marketing. Further, this change of belief could be achieved through social media marketing (Matz et al. 2017). The advantage would be that alternative finance peer-to-peer insurance companies would distinguish themselves from the incumbents through a positive narrative of risk-reduction through belonging to an in-group, rather than from having a contract with a big and perceived trustworthy company. Further research and FinTech start-ups could focus on identifying optimal matching groups based on mutual trust through social media data.

3.6 Transparency

The identification of optimal group matches and mutual trust could also help improving recruitment systems for alternative finance companies. For example, a big German software company cooperated with the Cambridge Psychometrics Centre on their recruitment. Through analysing Facebook data, interested persons could understand in seconds, which position or function would fit best to their personalities. The data was neither stored, nor shared with anyone. It was only used for informational purposes only. Such a system could be further developed towards personality matches within groups, and hardened against hidden biases in the data, which is more complex than one would expect. The example of Amazon's 'secret' internal recruitment system (Dastin 2018) shows that research on algorithmic ethics is crucial for every aspect of alternative finance that deals with psychometrics or machine learning. As an example, first names usually reflect the gender of their bearers, but they also indicate the country someone comes from, their religion (e.g., 'Christian' or 'Ali'), and ethnicity (e.g., 'Alfred Miller' or 'Liwei Yang'). However, 'Gideon Smith' may very well be an American of African

descent since slaves either were given or chose the surnames from their owner (Burnard 2001; Paterson 2001). Often, their first names were given by their owners, as well. Since religion provided slaves with a social measure to unite against injustice (Raboteau 1999), they very often bear Christian first names—even their descendants today. Also, based on housing prices, certain historically disadvantaged ethnicities may cluster-up in certain areas of cities. If an unsupervised algorithm determines creditworthiness based on personal or geographic data, societal injustice of the past may lead to biased decisions in the present that are not obvious—especially if control data is missing, for example about ethnicity.

But even if all bias is excluded on a theoretical level, the deployment of complex self-learning algorithms could create a situation, in which minorities still find themselves disadvantaged. For such situations, alternative finance companies need to ensure that their systems are transparent, based on former evidence and thus predictable, robust against manipulation, cannot get abused from bureaucrats to shirk responsibility, and bare from the ‘tendency to not make innocent victims scream with helpless frustration’ (Bostrom and Yudkowsky 2014, p.3). On a technical level, this can be ensured by using technologies like the one in the wolf-husky example (Ribeiro et al. 2016), or preferring generalised regressions, decisions trees, and Bayesian networks over neural networks or genetic algorithms (Hastie et al. 2009). On a logistical level, it could be helpful to collect more data where possible to better control against bias. The discussion about adverse impact in psychometric tests is much older than the same discussion in machine learning, and since it often involves hiring decisions, it is bound to local labour laws. For example, while collecting data on ethnicity is a standard recruitment procedure in the US to protect employees against discrimination, it would be unthinkable in countries like Germany, France, or Italy to ask for it.

European companies are afraid of abuse of these data points since they could be abused for future discrimination. However, US companies are afraid of litigation if they cannot bring a proof that either their selection measures do not discriminate or that the measures are so important for job success that they cannot do without, for example, numerical reasoning for accountants. That means as well, that European companies will encounter greater difficulty to control their algorithms against racial bias. Alternative finance companies should not only consider that for their global recruitment practices, but also for their products. Collecting the ‘wrong’ information may lead to reputational damage. Alternative finance companies would benefit from following the existing regulations in the industry with regard to which algorithms could

be used, as well as governmental obligations like GDPR or the UK Data Protection Act.

Last, alternative finance companies should think about their societal impact when they deploy machine learning or psychometrics. The more information gets digitised, the more discussion needs to take place whether we want to connect every piece of information from every citizen in one big dystopian database. If, for example, the whole world would go cashless via blockchain-based technology, who would prevent abolishment of privacy and the creation of a full dependency on a digital infrastructure? What would happen if there would be an electricity outage over a couple of hours or days? Who would prevent abuse through negative interest rates? The establishment of the GDPR and the current discussion about Articles 11 and 13 in the EU Copyright directive are already repercussions for the technology industry based on monopolies, data abuses, and societal and political ramifications of ongoing digitisation. At the same time, governments demand unlimited access to data and justify it with fighting terrorism or paedophilia. Companies, in turn, react to such requests with stronger encryption, moving principal offices, and only partial compliance. In the wake of the scandals around PRISM and Cambridge Analytica, it is in the interest of state actors and corporations alike to craft sensible regulation and policies that build trust with consumers. These should be tightly aligned with controlling against geographic and racial biases that may discriminate against individuals and prevent companies entering markets or launching products, for example for group lendings. Making coordinated, systematic and strategic use of 'soft' information from human behaviour and psychometrics through the capabilities of machine learning is in the best interest of every alternative finance company, and the authors conclude that appropriate deployment of both should be positively reflected on their balance sheets, as well.

3.7 Importance of Regulations

Two important TechFin market entries, one TechFin whistle blow, and one central bank proposal occurred in 2019, which will have far-reaching consequences for Alternative Finance, and which demand immediate and consequent regulations. Apple announced its own credit card in cooperation with Goldman Sachs (Apple 2019), Facebook announced its own digital currency, called 'Libra' (Libra 2019b), Mark Carney, the governor of the Bank of England, which carries out active research on risks and the regulation of digital currencies (Digital Currencies 2019), proposed that a global digital currency, which resembles Facebook's 'Libra', 'could replace the

U.S. dollar as the world's reserve currency' (Bambrough 2019), and former Google employee Zachary Vorhies blew the whistle on Google's algorithmic manipulation of both politics and private lives of citizens (Project Veritas 2019).

The issue of a credit card from Apple follows the argumentation of the authors that the strategy of TechFin companies is to use convenience to convince their customers into connecting a growing number of their life-spaces in their ecosystems. This convenience and connecting function comes at the cost of a less competitive product, which only pays 1% of the purchase in daily cash if the physical card is used, as compared to the Amazon store card, which is issued by JPMorgan Chase, yet offers a tight integration into the eco-system, unique convenience features, and relevant data insight for its customers (Apple 2019). While the issue of branded credit cards is not novel, the big difference is that the Apple credit card, is connected to an entire eco-system that the customer is not meant to leave. This gives Apple much more power over their customers, and provides them with unprecedented insight into the complete life-spaces of their customers. Despite their strong stance on privacy, the mere possession of this data and the ability to generate psychometric profiles that can be abused to manipulate their customers in any direction 'desired', gives them stronger power over their users than most governments have. This should clearly be very soon addressed by regulators.

Another factor to consider about the credit card from Apple is that it is issued by Goldman Sachs, a company that is in the institutional client and investment business (Goldman Sachs 2018), which reaches through Apple a new clientele of end-customers. Though counter-intuitively, this constellation is aligned with the analysis of the authors. As consumer-oriented banks and incumbents will need to cooperate with FinTech companies to battle TechFin companies at their own approach, so will business-oriented banks have to cooperate with TechFin companies, that become more and more serious factors in Alternative Finance, and soon maybe in the financial markets, as well. With the issue of the digital currency 'Libra', Facebook cooperates with a closed consortium of 28 global financial incumbents, TechFin and FinTech companies like MasterCard, Paypal, Stripe, Visa, Ebay, Lyft, Uber, Vodafone, coinbase, Andreessen Horowitz, and Women's World Banking. Though based on block-chain technology, 'Libra' is a centralised currency, which is in contrast to the decentralised architecture of bitcoin (Libra 2019a).

Though communicated as helping the poor and under-banked, the architecture of 'Libra' raises wide-spread concerns and calls for regulation. These concerns are raised around three major aspects of 'Libra'. First, the organisation of 'Libra' itself is the antithesis to a decentralised block-chain like bitcoin,

where every participant can validate payments. The speed and scale necessary for daily life can only be handled through a semi-centralised ledger, in which fewer and more powerful members of the 'Libra' consortium validate payments. All transactions would be recorded, and due to the closed circle of consortium members that includes the two dominant issuers of credit cards, Visa and Mastercard, a new central permission authority would be created, which means that the concept of decentralisation would be effectively removed (Murphy and Waters 2019). According to European Central Bank's board member Yves Mersch, such a set-up is 'extremely concerning' due to a lender of last resort and its sole responsibility towards shareholders (Koranyi et al. 2019). Second, the distribution channel Facebook is under public scrutiny due to the alleged Russian interference in the last US presidential election, its influence through its large global users base, the danger of sharing data with other Facebook services for purposes of creating personal profiles (Murphy and Waters 2019), unethical psychological experiments with its users (Kramer et al. 2014), and the fact that they deliver a compatible digital wallet for 'Libra', called 'Calibra'. 'Calibra' is a subsidiary of Facebook, led by Kevin Weil, the co-creator of 'Libra', former Vice President of Product at the Facebook company Instagram and Senior Vice President of Product at Twitter. The main advantages over all other 'Libra' wallets will be its integration with Facebook's global user base, and its proximity to Facebook's internal technological and product infrastructure. Through this, it could enable Facebook to control 'Libra' indirectly, and act as its bank, with which other TechFin could not compete (Statt 2019). This aggregation of power in one organisation is problematic (Murphy and Waters 2019), especially since Facebook displayed in the past a lack of responsibility towards respecting the privacy and trust of its users, for example with the decision to make use of data from WhatsApp (Weissman and Weissman 2016). Facebook's past behaviour, technological capabilities, and global user base also worries the European Central Bank as 'Libra' is issued by '...the very same people who had to explain themselves in front of legislators in the United States and the European Union on the threats to our democracies resulting from their handling of personal data on their social media platform' (Koranyi et al. 2019). Therefore third, expected systemic risks range from the interference of 'Libra' with policing the financial system through central banks, over limiting the ability of governments to restrict monetary flows and sanctions, to undermining their sovereignty on currencies through the establishment of a parallel currency in times of crises. Therefore, the 'threat to the existing banks is severe', especially since the plan behind 'Libra' is to establish a currency that lasts 'hundreds of years' (Murphy and Waters 2019), and that

would be connected to a consortium of global private companies that were not elected democratically, and that are ultimately only responsible to their shareholders and not to nation states or political regions like the European Union. Depending on its usage and the amount of Euro held as a reserve, it even could negatively influence the control of the European Central Bank over its own currency, and therewith the 'monetary policy transmission mechanism by affecting the liquidity position of euro area banks, and undermine the single currency's international role' (!), wherefore Europe should not be 'tempted' by the 'treacherous promises of Facebook's siren call', according to European Central Bank's board member Yves Mersch (Koranyi et al. 2019).

These strong reactions from high-ranking officials represent very well the potential influence that Facebook's currency might have on the world financial markets and its institutions. Therefore it is crucial that 'European regulatory and supervisory authorities assert jurisdiction over Libra' in cooperation with their global counterparts to manage potential risks (Koranyi et al. 2019). US authorities struggle to find a regulatory pattern in which 'Libra' would fit. A hearing in the Senate Banking Committee raises comparable concerns as the European Union, and one Senator even calls Facebook a 'toddler' that plays with 'matchers' (*sic!*) and 'calls every arson a learning experience' (Harper 2019a). In a follow-up hearing to that in the House of Representatives Financial Services Committee, it was impossible to determine what 'Libra' exactly is, and how it could be classified for regulation, since it's a reserve-backed digital currency, but not a crypto-currency. However, consensus is that 'there is a tremendous anticompetitive component' to it, given the closed consortium behind it that controls all of its aspects (Harper 2019b). In the Senate hearing, Microsoft's 'embrace, extend, extinguish' strategy (Economist 2000) is mentioned, that Facebook has used it before, and that 'Libra' could be another manifestation of that (Harper 2019a). The courts found Microsoft guilty of illegal monopolistic behaviour (United States of America 2001), which indicates that 'Libra' in its current form could very well be controlled through regulating against its core features. For example, the closed consortium could be opened, the connection of Facebook or any of its companies or affiliates to it could be forbidden, the free holding of money in 'Libra' could be priced, the cheap money transfer around the world could be taxed, its framework functionality against consortium members like financial service providers Visa and Mastercard or technology companies like Lyft and Uber could be forbidden, 'Calibra' could be sold to banks or forbidden to offer credit cards, which they think of doing in the future (Statt 2019), and, most importantly, its global spread could be mitigated through giving the central banking and governmental financial control

and regulation authorities a premium seat in the consortium with a full right to veto each. Given the strong public reaction, the threat of law suits and regulations, and assumptions that the existence of a consortium with various members is just a ruse to hide Facebook's central role in 'Libra' (Statt 2019), Facebook already hired a lobbying firm to ease regulatory pressure on 'Libra' (Huillet 2019) and extended its public relations work and attendance in hearings (Tamny 2019). And as the Microsoft lawsuits were appealed and settled, the authors predict that a similar development will take place with 'Libra', 'Calibra' for Facebook. Future research should focus on the effect of technological advancements on the development of regulatory authorities, with a special focus on how this influences their cooperation.

While there is the argument that Libra is the reaction to long lasting currency instabilities and lack of trust towards authorities (Tamny 2019), 'Libra' rather establishes new forms of corporate control from TechFin companies between banks and citizens. However, it may impose at the same time a new discipline in the banking system, especially since it's tied to fiat currencies and needs real currencies for price stability. This opens the question for responsibility, for example, when the reserves of Libra won't cover the withdrawals and central banks don't want to support their new competitors, since the approach of 'Libra' is to replace global central banking-based financial system. If it would be recognised as a sovereign currency, Facebook could become a new central or 'shadow-bank' of an alternative financial system (Brown 2019). This risk could, in turn, start an innovative cycle within the central banking system, which may explain why the governor of the Bank of England, Mark Carney, proposes a replacement of the US dollar as world reserve currency with a new global digital currency similar to 'Libra', which he calls a 'Synthetic Hegemonic Currency', instead of allowing other currencies like China's renminbi to act as a replacement (Swint 2019). He further proposes stronger interactions of central banks with the technology companies. For example, these could leave funds overnight with central banks to have alternatives to private banks and thus lower costs. This would bring central banks in competition to private banks and a step closer to issuing digital currencies. This adaptation is important, since technology companies could just change the location of their payment plans, financial authorities need to adapt (Murphy and Waters 2019), especially since blockchain is an 'unstoppable force' (Harper 2019b). This proposal of Mark Carney is remarkable as that it indicates political, economic, and social changes on a global scale, and the development of a multi-polar, less Western-centric world. However, more corporate control and the weakening of financial authorities to a point that they need to adapt and cooperate more strongly, also imposes

a risk on the overall systemic stability. Therefore, it doesn't wonder citizen rights groups criticise 'Libra' and its ramification on the financial system as distributed oppression and a Trojan horse for surveillance (LibreNotLibra 2019c).

Since the election of Donald J. Trump as 45th President of the United States, discussion about elections interference from Russia or commercial actors like Cambridge Analytics emerged. Due to these discussions, social media and technology companies started deplatforming extremists from their services, and joined efforts to fight fake news. While these efforts are necessary, their own political bias sees only little discussion. In the Senate Committee on the Judiciary, Subcommittee on the Constitution hearing entitled 'Google and Censorship through Search Engines' the Senior Research Psychologist of American Institute for Behavioural Research and Technology, Robert Epstein, gave a testimony on election meddling due to changed search results from Google. Therein, he analysed the interference of Google with the election and estimated between 2.6 and 10.4 voters may have been given to Hillary Clinton in the 2016 presidential campaign, and that in the 2020 elections, up to 15 million votes may be reached (Epstein 2019). This interference is conducted by a private company that nobody elected, and which threatens the fabric of democracy through its actions. This described interference through search results resembles the techniques used by Cambridge Analytica (Cadwalladr and Graham-Harrison 2018). In the last year, Google experienced various whistle blowers that informed about the internal political bias in the organisation and its meddling with public opinions and elections. The most prominent of these is Zachary Vorhies, who informed the public through the alternative media journalists from Project Veritas about the way Google manipulates search results for political purposes. He published several internal and confidential documents from Google, that exposed biased censorship, election tampering, fake news, hiring practices, leadership trainings, machine learning algorithms, partisanship, and secret psychological research. In the latter category, an internal research paper called 'Algorithmic Discrimination from an Environmental Psychology Perspective: Stress-Inducing go/discrimination-and-stress Differential Treatment' stands out. It describes a stream of research to better understand the stress effects that algorithmic discrimination has on human beings, for example which forms of stressors are easy and which are hard to tolerate, how humans try to control it, how this can be mitigated through design principles, the effect the environment has on it, and how mobile ads can act as a stress relief (Project Veritas 2019). This form of research bears results that could easily be used to manipulate people into a desired voting behaviour, especially if their individual

psychometric properties and stress coping mechanisms are known. And, it overlaps with the direction in Robert Epstein's testimony. Future research should focus on better understanding how psychometrics may have been used by Google to assert influence on elections, and how this may affect large scale human behaviour as it is reflected in the global financial markets.

The above described derailments of our post-reality (Holmes 2016) and post-factual (Schwartz 2016) politics are part of a greater development towards surveillance capitalism, as phrased by Harvard scholar Shoshana Zuboff. It is based on global behavioural modification, which is used to modify and program human nature, as industrial capitalism influenced the natural world. As the industrial revolution found its destructive climax in fascism as an amalgam of society, economy, politics and military under one dominant rule, surveillance capitalism seeks to merge nation-states into transnational corporatocracies. And as citizens in fascism must behave like industrial cogwheels, the behaviours of citizens in post-truth corporatocracies are controlled by employment dependability and consumerism, which are maximised by means of surveillance (Zuboff 2019). And this change must not come like the Chinese social credit system as a law, but it rather comes through the systematic transformation of social structures by means of corporate efficiency and measurement thinking. As in corporations, good behaviours and connections with the right people are rewarded, and bad behaviours and connections with the wrong people are punished. In the core of such a system is not external force like in fascism, but internal control through perceived self-improvement that enables individuals to become 'better' citizens with more rights to consume, for example through a selfish ledger (Savov 2018). To achieve this, privacy, borders, and anonymous cash need to be overcome, and more data needs to be collected. In this development, TechFin companies outcompete incumbents through much lower transaction costs and well-established digital services that already reach billions of people (Murphy and Waters 2019). Since banking and technology merge, outdated banking regulations and authority practices cannot be the norm any longer, as these are as vulnerable to TechFin strategies as incumbents are. This is visible in the Senate hearings of Facebook on 'Libra' or the Cambridge Analytica scandal that did not result in any regulations yet, but only yielded more questions and unveiled the technological divide between unsavvy rulers and their adept subjects. Therefore, the authors conclude that as incumbents, regulators should partner-up with FinTech companies as well, for being able to oppose TechFin tactics. For example, companies like Telegram or Brave created their own digital cryptocurrencies that are private in nature, decentralised, and meant to empower the

individuals that use these. In an unlikely alliance, the knowledge on their development and implementation could benefit regulators to better stabilise currencies and countries.

Furthermore, Europe's GDPR could be the blueprint for future global financial regulations, as it was successful to control surveillance capitalism. For example, in Europe, WhatsApp user data will not be shared with its parent company Facebook (Gartenberg 2018) due to the GDPR. Also, the Achilles heel of all TechFin companies is that they are either a publisher or a platform. As a platform, Facebook would not be allowed to censor. As a publisher, it would be responsible for every comment. Lawmakers and regulators can easily force it into one of the categories, which would mean the end of their operations as they are today. Through sensible regulations and anti-trust laws, the most influential business practices can be controlled, as well: collection of user data and the connection of life-spaces. The easiest way ahead for this would be to break up these TechFin companies in their components as it has been done successfully with AT&T in 1982 (Naylor 2019). There is no reason why this shouldn't be possible with Google, Facebook, Amazon, and Apple, and 'Libra' could be a well-suited precedent for further regulatory actions against TechFin. On a side-note, in the science-fiction movie 'Equilibrium', a dystopian and totalitarian government is described that is based on surveillance, shaping the behaviours of individuals through propaganda, and controlling their emotions; its capitol is called 'Libra' (Lawrence 2017). The authors propose to conduct further research on how the breakup of FinTech companies could be beneficial for democracy, societal stability, and financial market stability. This research should put a strong emphasis on how a breakup could happen alongside life-spaces instead of entities. Additional research should be conducted on how the GDPR could form a blueprint for financial regulations. Finally, the authors propose additional research on how the guiding principles identified in this chapter could benefit this process: trust, acceptance, privacy, together with freedom of choice, reciprocity, control, security, personalisation, advanced analytics, and portability (World Economic Forum 2018). Also, extensive research should be conducted on how the absence of data could be used by algorithms to infer predictions on individuals, as well. For example, following the GDPR, persons have the right to be forgotten. However, making use of this right, becomes another predictor, maybe for less favourable outcomes. Therefore, the right of having the act of being forgotten to be forgotten, as well, should be explored. Additional research should be conducted on the influence of network data on predicting the behaviour of individuals. Bagrow et al. (2019) showed that 95% of individual behaviour can be predicted from the network

of a person. Therefore, regulation should extend over the network of a person or potential persons as shadow profiles, as well.

4 Discussion and Outlook

Psychometrics permeates most aspects of the fourth industrial revolution since it enables a closer interaction of artificial with human intelligence and their subsequent collective behaviour. In the not too distant future, humans become a bit more machine with new input and output methods through brain links, while machines become more human through advanced neural and genetic algorithms. However, on the way to there, all industries go through fundamental transformational processes. In this article, the authors covered first new sources and use cases of data. Second, the deployment of machine learning and artificial intelligence over and above classical methodology from economics and econometrics. And, third, the accompanying ethical risks, regulations, and opportunities.

These transformational processes of all industries are in the centre of the discussion around tapping into new sources and finding new use cases for data. In the finance industry, incumbents experience structural and regulatory barriers, challenger corporations create and defend niches, while small FinTech start-ups change customer expectations through disruptive innovation. Their unique competitive edge is the deployment of new technologies, flexibility, and speed of adaptation. This development forces incumbents to adapt, since they now compete for the same spot as challenger corporations and emerging companies. However, they can make use of People Analytics to quantify people decisions with strategic reorganisations and tactical imperatives, and also use their People Analytics capabilities for product development. Also, they can make use of internal, C-Level sponsored innovation incubators to keep up with the new innovation pace of the industry. These technological changes invoke societal changes and vice versa. Global connectivity, omnipresence of ambient computing, and the emergence of close-to-brain computing will result in much more collectivist societies. The central hypothesis of the authors is that only the adaptation of both societal and technological changes will ensure business continuity.

The deployment of machine learning and artificial intelligence is necessary due to restrictions of econometrics and theory-driven approaches. First, existing approaches will be augmented by machine learning and second, both will be gradually replaced by deep learning. This development demands strong sense of responsibility, since most modern algorithms are black boxes,

and research only recently found ways to make them more transparent. The main difference between econometric techniques and machine learning is the inference of causality through theory and experimental setup in the former versus the focus on predictions in the latter. Based on the growing amount of data and lack of sufficient theory, predictions have a stronger influence already. Counter-intuitively, the authors recommend using less data and focusing rather on which data can be ignored than which data can be used. Existing data must be better cleaned, connected, standardised, and exploited fully.

Discussing ethical risks, regulations, and opportunities, the authors conclude that trust, acceptance, privacy, freedom of choice, reciprocity, and transparency are necessary for business continuity and even for defending the competitive edge. As products and offerings become ever more tailored to individuals, customers become increasingly selective with the companies they interact with. Therefore, alternative finance should consider their clients as part of the team, and develop new ways of personal interaction, based on trust. This means that companies need to understand psychology very well. But this means as well, that customers increasingly influence internal decisions of companies. Based on strategic considerations, market opportunities, and psychological analysis, the authors conclude that the insurance industry may be the best target for disruption. However, the narrative should rather focus on social backing than on costs saving costs. Also, the marketing before product launch must be extended to form prior beliefs. This is in stark contrast to the prevailing mindset of even the most innovative FinTech start-ups.

Discussion and hints for future research are given along the chapter to represent the complexity and importance of each topic. Beyond that, the authors contribute strategic considerations and predictions for future developments based on identified trends. These are by no way validated and should not be taken as an investment or legal hint, but as a discussion for academic exchange. It has to be stated that between the writing of this chapter and the publication of the book several years passed, in which quite a few changes happened, though all in line with the predictions of the authors. The authors refrained from re-writing the chapter based on the latest developments, mainly around the finalisation of the OpenSchufa project, the recent presidential elections in the USA, and the rebranding and downsizing of 'Libra' into 'diem', since the goal was to derive strategic information and academic insights based on ongoing general trends, and not to represent the latest state of affairs. A follow-up publication may or may not take up

these changes, however the authors consider it as not relevant to the extracted content and academic as societal debate that is to follow.

Finally, the authors wish to stress that deploying more psychometrics demands higher levels of empathy due to closer interaction of humans with the help of technology. Alternative finance has to think outside the box and beyond the panopticon. The readers should ask themselves whether they should do everything they are able to do. The movie 'Minority Report' was a surprisingly exact prediction of the shape of things to come. Which world they want to create for the next generation? Artificial intelligence is for the people today what radium was for the people a hundred years ago. Before knowing about the disastrous ramifications, people used radium for everything, including ingestion. The radium girls were one of its first victims, and their terrible sufferings should warn us today to fully embrace the future, but to fully embrace responsibility, as well.

5 Data Collected Through Today's Devices

The following table provides an overview of data that's collected with today's devices, and is in no way complete. It shall only serve as an example of what's already possible to detect, and what psychometricians can deduce from the data points (Table 1).

Table 1 Selection of ambient data with sources and use cases

Category	Origin	Data	Exemplary use
Behavioural data	smart phones and tablets	Location	Geographic positioning
		Magnetic sensors	Positioning
		Text	Sentiment analysis, personality
		Heat sensors	Health measure
		Acceleration	Sport and health measure
		Camera	Measure emotions
		Microphone	Voice analysis
		Air pressure sensor	Health and movement measure
		Passwords	Convenience feature
		Time of activity	Health measure
	Smart watches and fitness bands	Sleep data	Health measure
		Bluetooth and Wifi ID	Assistive devices
		Location	Movement measure
		Heart sensor	Stress measure
		Skin resistance	Stress measure, lie detector
		Text	Personality, sentiment analysis
		Camera	Emotion, lie detector
		Microphone	Emotion, lie detector
		Heat sensor	Emotional arousal, fear
		Time of activity	Health measure
Entry system	Sleep data	Health measure	
	Bluetooth and Wifi ID time in and out	Assistive devices Time tracking, social network graph	

(continued)

Table 1 (continued)

Category	Origin	Data	Exemplary use
		location	Social network graph
		sentiment	Attrition and motivation
		others in environment	Social network
		Location	Buying habits
	Credit and loyalty card	Product history	Buying habit, personality, desires, values
		Location history	Wealth, social network analysis, shopping behaviour
		Location	Work and living location
	Web browsers	Web history	Personality, politics, religion, values, gender, ethnicity, age
		Camera	Emotion, arousal, lie detector
		Microphone	Emotion, arousal, lie detector
		Passwords	Convenience
		Time of activity	Sleeping habits, work
		Typing behaviour	Identification
		Location	Work and home
	Computers and cloud services	Contacts and friends	Social network
		Likes and dislikes	Personality, values, desires, hopes, fears
		All communication data	Personality, intelligence, personality, values, desire, hopes, fear, attrition, motivation

Category	Origin	Data	Exemplary use
		Documents	Work style, personality, diligence
		Pictures and videos	Emotion, lie detector, age, gender, race, ethnicity, political and sexual orientation
		Financial data	Buying habits
		Location	Work and living
		Voice data	Emotion and arousal
		Passwords	Convenience
		Internet usage	Habits, interests, likes
		Time of activity	Sleeping habits
		Bluetooth and Wifi ID	Assistive devices
		Location	Working habits, fitness
	IoT devices and personal assistants	Contacts and friends	Social network, health, wealth, political ideas
		Social network data	Religion, sexual orientation
		Likes and dislikes	Potential performance, career path
		All communication data	Personality, lie detection
		Documents	Diligence, intelligence
		Pictures and videos	Ethnicity, hobbies
		Financial data	Ethnicity
		Location	Wealth
		Voice data	Age
		Passwords	Convenience
		Internet usage	Likes and dislikes

(continued)

Table 1 (continued)

Category	Origin	Data	Exemplary use
		Time of activity	Health, intelligence, personality
		Likes and dislikes	Personality
		Musical taste	Tendency to fraud
		Interests	Values
		Sleep data	Drug abuse
		Health data	Sexual activity
		Political opinion	Brand preference
	Public available data	Satellite imagery	Neighbourhood wealth
		Criminal records	Personality
		Surveillance and web cams	Complete behavioural profile
		Statistical data (e.g., Statistisches Bundesamt, Eurostat, ECB Statistical Data Warehouse, UN Data)	Control data
		Commercial data (e.g. Bloomberg, S&P)	Control data
		Police and homeland security reports	Deviant behaviour
Biometric Data	Facial/ body images	Face structure	Intelligence, sexual orientation
		Iris scan	Degree of felt pain
		Iris reaction	Lie
		Skin change	Lie
		Age	Motivation
		Gender	Values
		Ethnicity	Religious beliefs
		Emotional state	Willingness to buy
		Sexual orientation	Political opinion
		Intelligence	Education

Category	Origin	Data	Exemplary use
		3D face image	Intelligence
		Recognise body without face	Identification
		Recognise person on walk	Identification
	Fingerprint scans	Surface and deep level fingerprint	Identification
		Blood flow	Life/ death recognition
		Heat	Emotional state
	FBrain data	Electric activity	Level of concentration
		Electric stimulation	Level of concentration
		Opto-electric stimulation	Brain activity
		Emotion	Personal importance
		Personality	Profession
		Intelligence	Education
		Fingerprinting	Unique identification
		Content	Personality and values
		Social influence	Personality
		Language features	Willingness to kill
		Identification	Identification
	genetic data	Diagnosis	Potential diseases
		Location	Social embedding
	Facebook	Text data	Personality
		Employment data	Conscientiousness
		Education data	Intelligence
		Romantic partners	Socio-economic status
		List of family and friends	Life expectancy

(continued)

Table 1 (continued)

Category	Origin	Data	Exemplary use
		Political opinion	Personality
		Religious opinion	Education
		Group membership	Personality
		Interests and hobbies	Political ideas
		Pictures and Videos	Intelligence
		Voice Samples	Emotional state
		Products acquired and sold	Values, political opinion
		Linked Log-on data from applications	All data from these services that these share
	Twitter	Location	Identification
		List of friends and followers	Political opinion
		Text data	Personality
		Political opinion	Income
		Group membership	Interests
		Pictures and Videos	Sexual orientation
		Linked Log-on data from applications	All data these applications share
		Employment data	Fit for a certain role
	LinkedIn	Work samples	Personality
		Text data	Interests
		List of colleagues and friends	Risk behaviour
		Recommendations	Values and fears
		Linked social media data	Interests
		Linked Log-on data from applications	Music taste
		Data from Microsoft Office	Personality
		Political opinion	Brand preference
		Group membership	Personality

Category	Origin	Data	Exemplary use
	Patent applications and white papers	Pictures and Videos Text data	Sexual orientation Intelligence
	GitHub	Work sample Cooperators Text data Work sample Cooperators Employment data Friends and co-workers Personal opinion Text data	Conscientiousness Agreeableness Personality Intelligence Values Personality Wealth Sexual preference Openness to new experiences
	Conference attendance and employer homepages	Work samples Cooperators Employment data Video data Text data Text data Work samples Cooperators Employment data Video data	Extraversion Influence Conscientiousness Lie detection Personality Intelligence Diligence Extraversion Fit for job Sexual orientation
	news mentions		

Acknowledgements The authors wish to thank Timo Koch from Ludwigs Maximilian University of Munich for his outstanding contribution.

References

- Apple and Goldman Sachs launch their credit card. *The Economist*. ISSN 0013-0613.
- Schneider, Frank, and Sabrina Weber-Papen. 1996. *Psychiatrie, Psychosomatik und Psychotherapie ...in 5 Tagen*, 85. Heidelberg: Springer. ISBN 978-3-540-89049-2.
- United States of America, Appellee v. Microsoft Corporation, Appellant, 253 F.3d 34 (D.C. Cir. 2001). <https://law.justia.com/cases/federal/appellate-courts/F3/253/34/576095/>, 2001.
- Global Internet Usage. *Wikipedia*, November 2018. Page Version ID: 870834776.
- OpenSchufa: The first results. <https://okfn.de/blog/2018/11/openschufa-first-results/>, 2018.
- Annual-report-2018.pdf. Technical report, Goldman Sachs, 2019.
- Digital currencies. <http://www.bankofengland.co.uk/research/digital-currencies>, 2019.
- Google Document Dump---Project Veritas, 2019.
- Libra Association | A not-for-profit organization. <https://libra.org/en-US/association/>, 2019a.
- Libra White Paper | Blockchain, Association, Reserve. <https://libra.org/en-US/white-paper/>, 2019b.
- LibreNotLibra. <https://librenotlibra.info/>, 2019c.
- Abdelaziz, Salma. 2011. Report: Egyptian dad names child 'Facebook'. <http://www.cnn.com/2011/WORLD/meast/02/21/egypt.child.facebook/index.html>.
- SCHUFA AG. 2018. Stellungnahme zur aktuellen Medienberichterstattung. <https://www.schufa.de/de/ueber-uns/themenportal/detailseite/themenportal-detailseite.12544.jsp>.
- Amazon. 2018. What is Amazon Rekognition?---Amazon Rekognition. <https://docs.aws.amazon.com/rekognition/latest/dg/what-is.html>.
- Arora, Kim. 2017. Aadhaar may present 'interesting use cases' in the future: Facebook India MD---*Times of India*. <https://timesofindia.indiatimes.com/people/Aadhaar-may-present-interesting-use-cases-in-the-future-Facebook-India-MD/articleshow/55585783.cms>.
- Ash, S.E. 1951. Effects of group pressure on the modification and distortion of judgements. In *Groups: Leadership and Men*. New Brunswick, NJ: Rutgers University Press.
- Athey, Susan. 2018. The impact of machine learning on economics, p. 31.
- Awad, Edmond, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-Francois Bonnefon, and Iyad Rahwan. 2018. The

- moral machine experiment, 563 (7729): 59--64. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/s41586-018-0637-6>. <http://www.nature.com/articles/s41586-018-0637-6>.
- Badenhausen, Kurt. 2018. The world's most valuable brands 2018. <https://www.forbes.com/sites/kurtbadenhausen/2018/05/23/the-worlds-most-valuable-brands-2018/>.
- Bagrow, James P., Xipei Liu, and Lewis Mitchell. 2019. Information flow reveals prediction limits in online social activity. *Nature Human Behaviour* 3 (2, February): 122–128. ISSN 2397-3374. <https://doi.org/10.1038/s41562-018-0510-5>.
- Baker, Kat. 2010. Poor recruitment practices damage consumer brand and bottom line. <https://www.personneltoday.com/hr/poor-recruitment-practices-damage-consumer-brand-and-bottom-line/>.
- Balakrishnan, Matt Hunter. 2018. Anita. Apple earnings Q1 2018: How much money does Apple have? <https://www.cnbc.com/2018/02/01/apple-earnings-q1-2018-how-much-money-does-apple-have.html>.
- Bambrough, Billy. 2019. The Bank of England Governor proposed a radical competitor to Bitcoin and the U.S. dollar, August. <https://www.forbes.com/sites/billybambrough/2019/08/31/the-bank-of-england-governor-proposed-a-radical-competitor-to-bitcoin-and-the-us-dollar/>.
- Bartram, Dave. 2005. The great eight competencies: A criterion-centric approach to validation. *Journal of Applied Psychology* 90 (6): 1185–1203. ISSN 1939-1854, 0021-9010. <https://doi.org/10.1037/0021-9010.90.6.1185>.
- Berens, Brad. 2017. What comes after smartphones? | Center for the Digital Future. <https://www.digitalcenter.org/columns/what-comes-after-smart-phones/>.
- Bergen, Mark, and Jennifer Surane. 2018. Google and mastercard cut a secret ad deal to track retail sales. <https://www.bloomberg.com/news/articles/2018-08-30/google-and-mastercard-cut-a-secret-ad-deal-to-track-retail-sales>.
- Bjorkegren, Daniel, and Darrell Grissen. Behavior revealed in mobile phone usage predicts loan repayment. *SSRN Electronic Journal*. ISSN 1556-5068. <https://doi.org/10.2139/ssrn.2611775>.
- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012a. A 61-million-person experiment in social influence and political mobilization 489 (7415): 295–298. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/nature11421>. <http://www.nature.com/articles/nature11421>.
- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012b. A 61-million-person experiment in social influence and political mobilization 489 (7415): 295–298. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/nature11421>. <http://www.nature.com/articles/nature11421>.
- Bostrom, Nick, and Eliezer Yudkowsky. 2014. The ethics of artificial intelligence. In *The Cambridge handbook of artificial intelligence*, ed. Keith Frankish and William M. Ramsey, 316–334. Cambridge University Press. ISBN 978-1-139-04685-5.

- <https://doi.org/10.1017/CBO9781139046855.020>. https://www.cambridge.org/core/product/identifier/CBO9781139046855A027/type/book_part.
- Botsman, Rachel. 2017. Big data meets big brother as China moves to rate its citizens. ISSN 1357-0978. <https://www.wired.co.uk/article/chinese-government-social-credit-score-privacy-invasion>.
- Box, G.E.P., and G.M. Jenkins. 1976. *Time series analysis---Forecasting and control*. San Francisco: Holden-Day.
- Brandom, Russell. 2018. Even if you're not signed up, Facebook has a shadow profile for you. <https://www.theverge.com/2018/4/11/17225482/facebook-shadow-profile-files-zuckerberg-congress-data-privacy>.
- Brehm, Jack W. 1966. *A theory of psychological reactance*. New York: Academic Press.
- Breiman, Leo. 2001. *Random forests* 45: 5–32.
- Brown, Ellen. 2019. Facebook may pose a greater danger than Wall Street, June. <https://www.commondreams.org/views/2019/06/26/facebook-may-pose-greater-danger-wall-street>.
- Bunning, Jonny. 2018. Claps and cheers: Apple stores' carefully managed drama. ISSN 0261-3077. <https://www.theguardian.com/technology/2018/dec/03/claps-and-cheers-apple-stores-carefully-managed-drama>.
- Burnard, Trevor. 2001. lave naming patterns: Onomastics and the taxonomy of race in eighteenth-century Jamaica, 31 (3): 325–346. ISSN 0022-1953, 1530-9169. <https://doi.org/10.1162/002219500551550>. <http://www.mitpressjournals.org/doi/10.1162/002219500551550>.
- Cadwalladr, Carole, and Emma Graham-Harrison. 2018. Revealed: 50 million Facebook profiles harvested for Cambridge analytica in major data breach. ISSN 0261-3077. <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.
- Cameron, James E., Jennifer Voth, Susan B. Jaglal, Sara J.T. Guilcher, Gillian Hawker, and Nancy M. Salbach. "In this together": Social identification predicts health outcomes (via self-efficacy) in a chronic disease self-management program, 208: 172–179, 2018. ISSN 02779536. <https://doi.org/10.1016/j.socscimed.2018.03.007>. <https://linkinghub.elsevier.com/retrieve/pii/S027795361830100X>.
- Capgemini and LinkedIn. 2018. World Fintech Report.
- Carter, Shawn M. 2017. Mark Zuckerberg could be running for President in 2020. <https://www.cnbc.com/2017/08/15/mark-zuckerberg-could-be-running-for-president-in-2020.html>.
- Clark, Nicola. 2012. Selecting a Seatmate to make skies Friendlier. ISSN 0362-4331. <https://www.nytimes.com/2012/02/24/business/global/selecting-a-seatmate-to-make-skies-friendlier.html>.
- Cohn, Alain, Ernst Fehr, and Benedikt Herrmann. 2011. Social comparison in the workplace: Evidence from a field experiment, p. 36.
- Cole, Samantha. 2018. Sex workers say porn on Google Drive is suddenly disappearing. https://motherboard.vice.com/en_us/article/9kgwnp/porn-on-google-drive-error.

- Dastin, Jeffrey. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>.
- du Toit, Gerard, and Aaron Cheris. 2018. Banking's Amazon moment.
- du Toit, Gerard, Maureen Burns, Christy de Gooyer, David Phillips, Madhav Jalan, Lakshay Kataria, Sidhant Law, Akash Malhotra, Jocelyne Moyer, Pranav Singh, and Lavanya Srivastav. 2017. Evolving the customer experience in banking.
- Economist. 2000. Deadly embrace. *The Economist*. ISSN 0013-0613.
- Edwards, Jim. 2011. Social media is a tool of the CIA. Seriously. <https://www.cbsnews.com/news/social-media-is-a-tool-of-the-cia-seriously/>.
- Epstein, Robert. 2019. Google and censorship through search engines | United States Senate Committee on the Judiciary. <https://www.judiciary.senate.gov/meetings/google-and-censorship-through-search-engines>, 2019.
- European Commission. 2018. Are there restrictions on the use of automated decision-making? https://ec.europa.eu/info/law/law-topic/data-protection/ref-orm/rules-business-and-organisations/dealing-citizens/are-there-restrictions-use-automated-decision-making_en.
- Finnegan, Matthew. 2018. Microsoft brings AI transcription to OneDrive and SharePoint. <https://www.computerworld.com/article/3302160/office-365/microsoft-brings-ai-transcription-to-onedrive-and-sharepoint.html>.
- Fishman, Alex, and Douglas Hassebrock. 2010. A human driven data-centric approach to accountability: Analyzing data to prevent fraud, waste and abuse in stimulus spending: Gov 2.0 Expo 2010---O'Reilly Conferences, May 25--27, Washington, DC. <https://conferences.oreilly.com/gov2expo/gov2expo2010/public/schedule/detail/13996>.
- Forgas, Joseph P., Liz Goldenberg, and Christian Unkelbach. 2009. Can bad weather improve your memory? An unobtrusive field study of natural mood effects on real-life memory. *Journal of Experimental Social Psychology* 45 (1): 254–257. ISSN 00221031. <https://doi.org/10.1016/j.jesp.2008.08.014>.
- World Economic Forum. 2018. The appropriate use of customer data in financial services.
- International Monetary Fund. 2018. Report for selected countries and subjects. <https://www.imf.org/external/pubs/ft/weo/2018/01/weodata/weorept.aspx?sy=2017&ey=2017&scsm=1&ssd=1&sort=country&ds=.&br=1&c=512>
- Garris, Paul A., Michaux Kilpatrick, Melissa A. Bunin, Darren Michael, Q. David Walker, and R. Mark Wightman. 1999. Dissociation of dopamine release in the nucleus accumbens from intracranial self-stimulation, 398 (6722): 67–69. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/18019>. <http://www.nature.com/articles/18019>.
- Gartenberg, Chaim. 2018. WhatsApp won't share user data with Facebook in Europe, March. <https://www.theverge.com/2018/3/14/17120446/whatsapp-user-data-sharing-facebook-uk-privacy-ico-protection-gdpr-europe>.

- Geron, Tomio. 2013. Airbnb adds identity verification to end anonymity in sharing economy. <https://www.forbes.com/sites/tomiogeron/2013/04/30/airbnb-adds-identity-verification-in-big-step-for-sharing-economy/>.
- Girard, René, and James G. Williams. 1996. *The Girard reader*. Crossroad. ISBN 978-0-8245-1609-3 978-0-8245-1634-5.
- Glaeser, Edward L., Joseph Gyourko, and Albert Saiz. 2008. Housing supply and housing bubbles. *Journal of Urban Economics* 64 (2, September): 198–217. ISSN 00941190. <https://doi.org/10.1016/j.jue.2008.07.007>.
- Greenwald, Glenn, and Ewen MacAskill. 2013. NSA prism program taps in to user data of Apple, Google and others. ISSN 0261-3077. <https://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data>.
- Gu, Shihao, Bryan T. Kelly, and Dacheng Xiu. 2018. Empirical asset pricing via machine learning. ISSN 1556-5068. <https://doi.org/10.2139/ssrn.3159577>. <https://www.ssrn.com/abstract=3159577>.
- Guenole, Nigel, Sheri Feinzig, and Jonathan Ferrar. 2015. Employee privacy preferences in the world's major economies.
- Happ, Christian, André Melzer, and Georges Steffgen. 2016. Trick with treat--- Reciprocity increases the willingness to communicate personal data, 61: 372–377. ISSN 07475632. <https://doi.org/10.1016/j.chb.2016.03.026>. <https://linkinghub.elsevier.com/retrieve/pii/S0747563216301935>.
- Harper, Colin. 2019a. Bitcoin news: In senate hearing, Facebook's Iffy reputation looms over libra plans, July. <https://bitcoinmagazine.com/articles/senate-hearing-facebooks-iffy-reputation-looms-libra-plans>.
- Harper, Colin. 2019b. Libra in Congress: A chance to define what Bitcoin is and what Libra isn't, July. <https://bitcoinmagazine.com/articles/libra-congress-chance-define-bitcoin-libra-isnt>.
- Harwell, Drew. 2018. Your banking data was once off-limits to tech companies. Now they're racing to get it. <https://www.washingtonpost.com/technology/2018/08/07/your-banking-data-was-once-off-limits-tech-companies-now-theyre-racing-get-it/>.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The elements of statistical learning*, 2nd ed. Springer-Verlag.
- Heater, Brian. 2017. Google ends its hands free mobile payment pilot after less than a year, but promises more to come. <http://social.techcrunch.com/2017/02/02/google-hands-free/>.
- Holmes, Jack. 2016. Trump's campaign manager offered her most Brilliant defense yet of Trump's lies, September. <https://www.esquire.com/news-politics/news/a48906/trump-campaign-manager-lester-holt/>.
- Huillet, Marie. 2019. Facebook hires lobbying firm to ease regulatory pressure on Libra, August. <https://cointelegraph.com/news/facebook-hires-lobbying-firm-to-ease-regulatory-pressure-on-libra>.
- Hull, Clark L. 1952. *Clark L. Hull: A history of psychology in autobiography*. Worcester, MA: Clark University Press.
- IBM. 2017. 10 Key marketing trends for 2017.

- DNA India. 2018. What are Aadhaar-enabled ATMs and how will they help rural India. <https://www.dnaindia.com/business/report-what-are-aadhaar-enabled-atms-and-how-will-they-help-rural-india-2625071>.
- John, Leslie K., George Loewenstein, and Scott I. Rick. 2014. Cheating more for less: Upward social comparisons motivate the poorly compensated to cheat. *Organizational Behavior and Human Decision Processes* 123 (2, March): 101–109. ISSN 07495978. <https://doi.org/10.1016/j.obhdp.2013.08.002>.
- Judge, Timothy A., Carl J. Thoresen, Vladimir Pucik, and Theresa M. Welbourne. 1999. Managerial coping with organizational change: A dispositional perspective. p. 16.
- Kahneman, Daniel, and Amos Tversky. 1979. Prospect theory: An analysis of decisions under risk, 47 (2). <http://www.its.caltech.edu/~camerer/Ec101/ProspectTheory.pdf>.
- Kateliiev, Preslav. 2017. 4 apps that Apple banned only to “steal” their features with the next iOS update. https://www.phonearena.com/news/4-apps-that-apple-banned-only-to-steal-their-features-with-the-next-ios-update_id97171.
- Kaufman, Rob. 2016. 5 factors that determine a FICO score. <https://blog.myfico.com/5-factors-determine-fico-score/>.
- Kelion, Leo. 2016. Barclays joins Apple pay in UK. <https://www.bbc.com/news/technology-35968102>.
- Koranyi, Balazs, Peter Graff, and Louise Heavens. 2019. Europe should ignore ‘treacherous promises’ of Facebook’s Libra... *Reuters*, September.
- Kosinski, Michal, Sandra C. Matz, Samuel D. Gosling, Vesselin Popov, and David Stillwell. 2015. Facebook as a research tool for the social sciences: Opportunities, challenges, ethical considerations, and practical guidelines, 70 (6): 543–556. ISSN 1935-990X, 0003-066X. <https://doi.org/10.1037/a0039210>. <http://doi.apa.org/getdoi.cfm?doi=10.1037/a0039210>.
- Kramer, Adam D.I., Jamie E. Guillory, and Jeffrey T. 2014. Hancock. Experimental evidence of massive-scale emotional contagion through social networks, 11 (24): 8788–8790.
- Lawrence, Kathryn. 2017. TODAY’S DYSTOPIA: Equilibrium, September. <https://www.montag.wtf/todays-dystopia-equilibrium/>.
- Lewin, Kurt. 1939. Field theory and experiment in social psychology: Concepts and methods. *American Journal of Sociology* 44 (6): 868–896.
- Lewin, Kurt. 1943. Defining the ‘field at a given time’. *Psychological Review* 50 (3): 292–310.
- Lindstrom, Martin. 2008. How Apple, others have cultivated religious followings, p. 4.
- Lomas, Natasha. 2018. Yes Facebook is using your 2FA phone number to target you with ads. <http://social.techcrunch.com/2018/09/27/yes-facebook-is-using-your-2fa-phone-number-to-target-you-with-ads/>.

- MacDorman, Karl F., and Hiroshi Ishiguro. 2006. The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies* 7 (3): 297–337. ISSN 1572-0373. <https://doi.org/10.1075/is.7.3.03mac>. <https://www.jbe-platform.com/content/journals/10.1075/is.7.3.03mac>.
- Maslow, Abraham H. 1943. A theory of human motivation. *Psychological Review* 50: 370–396.
- Masnack, Mike. 2015. Green bubbles: How Apple quietly gets iPhone users to hate Android users. <https://www.techdirt.com/articles/20150211/05455029985/green-bubbles-how-apple-quietly-gets-iphone-users-to-hate-android-users.shtml>.
- Mason, Winter A., Frederica R. Conrey, and Eliot R. Smith. 2007. Situating social influence processes: Dynamic, multidirectional flows of influence within social networks, 11 (3): 279–300. ISSN 1088-8683, 1532-7957. <https://doi.org/10.1177/1088868307301032>. <http://journals.sagepub.com/doi/10.1177/1088868307301032>.
- Matz, S.C., M. Kosinski, G. Nave, and D. J. Stillwell. 2017. Psychological targeting as an effective approach to digital mass persuasion, 114 (48): 12714–12719. ISSN 0027-8424, 1091-6490. <https://doi.org/10.1073/pnas.1710966114>. <http://www.pnas.org/lookup/doi/10.1073/pnas.1710966114>.
- McFadden, Leslie. 2010. What makes up your credit score? <https://www.cnn.com/id/36737279>.
- McWaters, Jesse. 2015. The future of financial services how disruptive innovations are reshaping the way financial services are structured, provisioned and consumed.
- Nicas, Jack, and Cade Metz. 2018. Apple hires Google's AI chief. <https://www.nytimes.com/2018/04/03/business/apple-hires-googles-ai-chief.html?smid=tw-nytimesbits&smtyp=cur>.
- Microsoft. 2018a. Microsoft services agreement. <https://www.microsoft.com/en-us/servicesagreement/>.
- Microsoft. 2018b. Personal shopping assistant. <https://psa.microsoft.com/home/index.html>.
- Miller, Neal E. 1944. Experimental studies of conflict behavior. *Personality and Behavior Disorders: A Handbook Based on Experimental and Clinical Research*, 431–465.
- Miller, Neal E. 1951. Comments on theoretical models: Illustrated by the development of a theory of conflict behavior. *Journal of Personality*, 82-100.
- Mühlbauer, Peter. 2019. Google macht Banken Konkurrenz. <https://www.heise.de/tp/features/Google-macht-Banken-Konkurrenz-4266284.html>.
- Mullainathan, Sendhil, and Jann Spiess. 2017. Machine learning: An applied econometric approach, 31 (2): 87–106. ISSN 0895-3309. <https://doi.org/10.1257/jep.31.2.87>. URL <http://pubs.aeaweb.org/doi/10.1257/jep.31.2.87>.
- Murphy, Hannah, and Richard Waters. 2019. Facebook's full-frontal assault on finance, June. <https://www.ft.com/content/2b9e2190-940a-11e9-aea1-2b1d33ac3271>.

- Naylor, Brian. 2019. Could the old AT&T break-up offer lessons for big tech today?, June. <https://www.npr.org/2019/06/26/736344175/could-the-old-at-t-break-up-offer-lessons-for-big-tech-today>.
- Neate, Rupert. 2018. Over \$119bn wiped off Facebook's market cap after growth shock. ISSN 0261-3077. <https://www.theguardian.com/technology/2018/jul/26/facebook-market-cap-falls-109bn-dollars-after-growth-shock>.
- Nellis, Stephen, and Cate Cadell. 2018. Apple moves to store iCloud keys in China, raising human rights fears. <https://www.reuters.com/article/us-china-apple-icloud-insight-idUSKCN1G8060>.
- Ngo, Theresa. 2013. The Box-Jenkins methodology for time series models, p. 11.
- Nichols, Shaun. 2017. Apple: Sure, we banned VPN iOS apps in China, but, um, er, art!. https://www.theregister.co.uk/2017/11/22/apple_vpns_china_skype/.
- Nolan, Steve. 2013. Revealed: Google and Facebook DID allow NSA access to data and were in talks to set up 'spying rooms' despite denials by Zuckerberg and page over PRISM project. <https://www.dailymail.co.uk/news/article-2337863/PRISM-Google-Facebook-DID-allow-NSA-access-data-talks-set-spying-rooms-despite-denials-Zuckerberg-Page-controversial-project.html>.
- Obschonka, Martin, Michael Stuetzer, Peter J. Rentfrow, Neil Lee, Jeff Potter, and Samuel D. Gosling. 2018. Fear, Populism, and the Geopolitical Landscape: The " Sleeper Effect " of Neurotic Personality Traits on Regional Voting Behavior in the 2016 Brexit and Trump Elections, 9 (3): 285–298. ISSN 1948–5506, 1948–5514. <https://doi.org/10.1177/1948550618755874>. <http://journals.sagepub.com/doi/10.1177/1948550618755874>.
- O'Malley, James. 2015. Here's why Barclays doesn't support Apple Pay. <http://www.gizmodo.co.uk/2015/06/heres-why-barclays-doesnt-support-apple-pay/>.
- OpenSchufa. 2018. OpenSCHUFA---cracking the SCHUFA code. <https://www.starnext.com/en/openschufa>.
- Oreg, Shaul. 2003. Resistance to change: Developing an individual differences measure. 88 (4): 680–693. ISSN 1939-1854, 0021-9010. <https://doi.org/10.1037/0021-9010.88.4.680>. <http://doi.apa.org/getdoi.cfm?doi=10.1037/0021-9010.88.4.680>.
- Orlowski, Andrew. 2010. Facebook founder called trusting users dumb f*cks. https://www.theregister.co.uk/2010/05/14/facebook_trust_dumb/.
- Park, Gregory, H. Andrew Schwartz, Johannes C. Eichstaedt, Margaret L. Kern, Michal Kosinski, David J. Stillwell, Lyle H. Ungar, and Martin E. P. Seligman. 2015. Automatic personality assessment through social media language, 108 (6): 934–952. ISSN 1939-1315, 0022-3514. <https://doi.org/10.1037/pspp0000020>. <http://doi.apa.org/getdoi.cfm?doi=10.1037/pspp0000020>.
- Paterson, David. 2001. A perspective on indexing slaves' names, 64 (1): 132–142. ISSN 0360-9081. <https://doi.org/10.17723/aarc.64.1.th18g8t6282h4283>. <http://americanarchivist.org/doi/10.17723/aarc.64.1.th18g8t6282h4283>.
- Pavouir, Ben. 2017. What a Facebook experiment did to news in Cambodia. <https://www.bbc.com/news/world-asia-41801071>.

- Prabhu, Jaideep. 2017. Frugal innovation: Doing more with less for more, 375 (2095): 20160372. ISSN 1364-503X, 1471-2962. <https://doi.org/10.1098/rsta.2016.0372>. <http://rsta.royalsocietypublishing.org/lookup/doi/10.1098/rsta.2016.0372>.
- Quittner, Jeremy. 2012. Banks to use social media data for loans and pricing. <https://www.americanbanker.com/news/banks-to-use-social-media-data-for-loans-and-pricing>.
- Raboteau, Albert J. 1999. *African-American religion*, vol. 1. USA: Oxford University Press.
- Reinsel, David, John Gantz, and John Rydning. 2018. The digitization of the world from edge to core, p. 28.
- Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. 2016. “Why should I trust you?”: Explaining the predictions of any classifier. <http://arxiv.org/abs/1602.04938>.
- Rodotá, Stefano, and Rafael Capurro. 2005. Ethical aspects of ICT implants in the human body, 10 (1). ISSN 1613-1142, 1430-9017. <https://doi.org/10.1515/9783110182521.501>. <https://www.degruyter.com/view/j/jfwe.2005.10.issue-1/9783110182521.501/9783110182521.501.xml>.
- Roose, Kevin. 2018. How Facebook’s data sharing went from feature to bug. ISSN 0362-4331. <https://www.nytimes.com/2018/03/19/technology/facebook-data-sharing.html>.
- Rosenthal-von der Pütten, Astrid M., and Nicole C. Krämer. 2014. How design characteristics of robots determine evaluation and uncanny valley related responses, 36: 422–439. ISSN 07475632. <https://doi.org/10.1016/j.chb.2014.03.066>. <https://linkinghub.elsevier.com/retrieve/pii/S0747563214001940>.
- Safi, Michael. 2018. Indian court upholds legality of world’s largest biometric database. ISSN 0261-3077. <https://www.theguardian.com/world/2018/sep/26/indian-court-upholds-legality-of-worlds-largest-biometric-database>.
- Satish, Balla. 2018. How WhatsApp helped turn a village into a mob. <https://www.bbc.com/news/world-asia-india-44856910>.
- Savov, Vlad. 2018. Google’s Selfish Ledger is an unsettling vision of Silicon Valley social engineering. <https://www.theverge.com/2018/5/17/17344250/google-x-sel-fish-ledger-video-data-privacy>.
- Schiff, N.D., J.T. Giacino, K. Kalmar, J.D. Victor, K. Baker, M. Gerber, B. Fritz, B. Eisenberg, J. O’Connor, E.J. Kobylarz, S. Farris, A. Machado, C. McCagg, F. Plum, J.J. Fins, and A.R. Rezai. 2007. Behavioural improvements with thalamic stimulation after severe traumatic brain injury, 448 (7153): 600–603. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/nature06041>. <http://www.nature.com/articles/nature06041>.
- Schmitt, Neal. 2014. Personality and cognitive ability as predictors of effective performance at work, 1 (1): 45–65. <https://doi.org/10.1146/annurev-orgpsych-031413-091255>. <https://doi.org/10.1146/annurev-orgpsych-031413-091255>.
- Schneider, Benjamin. 1987. The people make the place. *Personnel psychology* 40 (3): 437–453.

- Schneider, Benjamin, and Dave Bartram. 2017. Aggregate personality and organizational competitive advantage. *Journal of Occupational and Organizational Psychology* 90 (4): 461–480. ISSN 09631798. <https://doi.org/10.1111/joop.12180>.
- SCHUFA. 2018a. Transparentes Scoreverfahren. <https://www.schufa.de/de/ueber-uns/daten-scoring/scoring/transparente-scoreverfahren/>.
- SCHUFA. 2018b. Was ist scoring? <https://www.schufa.de/de/ueber-uns/daten-scoring/scoring/scoring/>.
- SCHUFA. 2018c. Wie funktioniert Scoring bei der SCHUFA? <https://www.schufa.de/de/ueber-uns/daten-scoring/scoring/scoring-schufa/>.
- Schultz, Jeff. 2017. How much data is created on the Internet each day? | Micro Focus Blog. <https://blog.microfocus.com/how-much-data-is-created-on-the-internet-each-day/>.
- Schwab, Klaus. 2017. *The fourth industrial revolution*. Currency.
- Schwartz, Ian. 2016. George Will: “Post-factual politics” From campaign still exists, Nixon more of a statesman than current leadership, November. http://www.realclearpolitics.com/video/2016/11/28/george_will_post-factual_politics_from_campaign_nixon_statesman_trump.html.
- Selye, Hans. 1955. Stress and disease. *Science*, 625–631.
- Shah, Nimish. 2015. BIG DATA FOR BANKING SURVEY RESULTS, p. 12.
- Shakya, Holly B., and Nicholas A. Christakis. 2017. Association of Facebook use with compromised well being a longitudinal study, 185 (3): 203–211.
- Shepardson, David. 2018. Facebook confirms data sharing with Chinese companies. <https://www.reuters.com/article/us-facebook-privacy-congress-idUSKCN1J11TY>.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. Mastering the game of go without human knowledge. *Nature* 550 (7676, October): 354–359. ISSN 0028-0836, 1476-4687. <https://doi.org/10.1038/nature24270>.
- Smith, Chris. 2017. Apple might be helping governments spy on users in real time--BGR. <https://bgr.com/2017/06/06/apple-spying-on-iphone-users/>.
- Spence, Kenneth W., and Janet A. Taylor. 1951. Anxiety and strength of the ucs as determiners of the amount of eyelid conditioning. *Journal of Experimental Psychology* 42: 183–188.
- Spence, Kenneth W., Janet A. Taylor, and Rhoda Ketchel. 1956. Anxiety (drive) level and degree of competition in paired-associates learning. *Journal of Experimental Psychology* 52 (5): 306–310.
- Statt, Nick. 2019. Facebook’s Calibra is a secret weapon for monetizing its new cryptocurrency, June. <https://www.theverge.com/2019/6/18/18682838/facebook-digital-wallet-calibra-libra-cryptocurrency-kevin-weil-david-marcus-interview>.
- Stoycheff, Elizabeth. 2016. Under surveillance: Examining Facebook’s spiral of silence effects in the wake of NSA Internet monitoring, 93 (2): 296–311. ISSN

- 1077-6990, 2161-430X. <https://doi.org/10.1177/1077699016630255>. <http://journals.sagepub.com/doi/10.1177/1077699016630255>.
- Swint, Brian. 2019. Carney urges Libra-like reserve currency to end dollar dominance. *Bloomberg.com*, August.
- Tamny, John. 2019. Facebook's Libra is a global reaction to endless currency instability, August. <https://www.forbes.com/sites/johntamny/2019/08/04/with-her-haughty-dismissal-of-facebooks-libra-rana-foroohar-unwittingly-exposed-her-own-naivete/>.
- Teivainen, Aleks. 2018. 300,000 Finns have stopped using Facebook since April, says social media blogger. <http://www.helsinkitimes.fi/finland/finland-news/domestic/15782-300-000-finns-have-stopped-using-facebook-since-april-says-social-media-blogger.html>.
- Tellis, Gerard J., Stefan Stremersch, and Eden Yin. 2003. *The international takeoff of new products: The role of economics, culture, and country innovativeness*. 22 (2): 188–208.
- Toesland, Finbarr. 2017. How banks use data. <https://www.raconteur.net/finance/how-banks-use-data>.
- Trout, Jonah, Chris Copenhaver, Hilary Halpern, Shuang Hao, and Vincent Tran. 2016. The complex regulatory landscape for FinTech an uncertain future for small and medium-sized enterprise lending.
- van der Linden, Sander. 2017. The nature of viral altruism and how to make it stick, 1 (3): 0041. ISSN 2397-3374. <https://doi.org/10.1038/s41562-016-0041>. <http://www.nature.com/articles/s41562-016-0041>.
- Van Der Linden, Sander. 2018. The future of behavioral insights: On the importance of socially situated nudges, pp. 1–11. ISSN 2398-063X, 2398-0648. <https://doi.org/10.1017/bpp.2018.22>. https://www.cambridge.org/core/product/identifier/S2398063X18000222/type/journal_article.
- VanVugt, Mark. 2008. Gender differences in cooperation and competition: The male warrior hypothesis. <http://doi.apa.org/get-pe-doi.cfm?doi=10.1037/e554382012-018>.
- Varian, Hal R. 2014. Big Data: New tricks for econometrics, 28 (2): 3–28. ISSN 0895-3309. <https://doi.org/10.1257/jep.28.2.3>. <http://pubs.aeaweb.org/doi/10.1257/jep.28.2.3>.
- Vergleich.de. 2018. Die SCHUFA: Was Bonität und SCHUFA-Score beeinflusst. <https://www.vergleich.de/schufa.html>.
- Vincent, James. 2016. Twitter taught Microsoft's friendly AI chatbot to be a racist asshole in less than a day. <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>.
- Vincent, James. 2018. Google 'fixed' its racist algorithm by removing gorillas from its image-labeling tech. <https://www.theverge.com/2018/1/12/16882408/google-racist-gorillas-photo-recognition-algorithm-ai>.
- Walsh, Toby. 2015. Open letter on autonomous weapons. <https://futureoflife.org/open-letter-autonomous-weapons/>.

- Wan, Feng, Peter J. Williamson, and Eden Yin. 2015. Antecedents and implications of disruptive innovation: Evidence from China, 39-40: 94–104. ISSN 01664972. <https://doi.org/10.1016/j.technovation.2014.05.012>. <http://linkinghub.elsevier.com/retrieve/pii/S0166497214000820>.
- Wei, Xingjie, and David Stillwell. 2017. How smart does your profile image look? Estimating intelligence from social network profile images.
- Weisbaum, Herb. 2018. Zuckerberg's apology tour has not done much to regain user trust. <https://www.nbcnews.com/business/consumer/trust-facebook-has-dropped-51-percent-cambridge-analytica-scandal-n867011>.
- Weissman, Cale Guthrie, and Cale Guthrie Weissman. 2016. WhatsApp breaks promise, will now share user data with Facebook, August. <https://www.fastcompany.com/4017734/whatsapp-breaks-promise-will-now-share-user-data-with-facebook>.
- Welch, Chris. 2018. Google just gave a stunning demo of Assistant making an actual phone call. <https://www.theverge.com/2018/5/8/17332070/google-assistant-makes-phone-call-demo-duplex-io-2018>.
- White, Mathew P., Sabine Pahl, Marc Buehner, and Andres Haye. 2003. Trust in risky messages: The role of prior attitudes, 23 (4): 717–726. ISSN 0272-4332, 1539-6924. <https://doi.org/10.1111/1539-6924.00350>. <http://doi.wiley.com/10.1111/1539-6924.00350>.
- Wilkens, Andreas. 2012. Bericht: Schufa will Daten in sozialen Netzwerken nutzen. <https://www.heise.de/newsticker/meldung/Bericht-Schufa-will-Daten-in-sozialen-Netzwerken-nutzen-1612450.html>.
- Willke, Helmut. 2000. *Systemtheorie I*. UTB: Grundlagen.
- Yoo, Tony. 2016. Your iPhone is silently uploading your metadata to Apple's iCloud servers. <https://www.businessinsider.com.au/your-iphone-is-silently-uploading-your-metadata-to-apples-servers-2016-11>.
- Zuboff, Shoshana. 2019. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.



Neural Networks and Deep Learning: A Paradigm Shift in Information Processing, Machine Learning, and Artificial Intelligence

Stephen Fitz and Peter Romero

1 Introduction

The sub-field of AI focusing on deep (multi-layered) neural architectures and the associated algorithms is collectively known as Deep Learning. The success of Deep Learning has been so great in recent years, that all of modern AI is basically the study of deep neural architectures. Deep Learning models hold state-of-the-art results on virtually all AI tasks, and new discoveries are made daily. In many real-world Machine Learning tasks, in particular, those with perceptual input, such as Vision and Speech, the mapping from raw data to the output is often a complicated function with many factors of variation. Previously, to achieve decent performance on such tasks, significant effort had to be put to engineer hand-crafted features. Deep Learning algorithms aim to learn feature hierarchies with features at higher levels in the hierarchy formed by the composition of lower level features. This automatic learning of features has been demonstrated to uncover underlying structure in the data (cf. the manifold hypothesis in data science). The first success came in the fields of

S. Fitz

Keio University Global Research Institute, Tokyo, Japan

e-mail: stephenf@keio.edu

P. Romero (✉)

University of Cambridge, Cambridge, UK

e-mail: [peterr@cam.ac.uk](mailto:peter@cam.ac.uk)

vision and speech processing. Second wave of deep learning advances brought significant contributions to broader field of NLP, and reinforcement learning. Unsupervised learning and applications to social sciences are currently on the rise. This chapter aims to introduce the basic building blocks of modern neural Artificial Intelligence (AI) systems to alternative finance audience, and other nonexperts in the field of AI. Deep Learning architectures composed of those components are being applied in a wide range of domains. In particular gated recurrent cells such as LSTM (which we will study later in the chapter) are especially useful in finance due to their ability to extract signal over extended periods in noisy, high-dimensional time series data. At the end of the chapter we'll examine how combining several ideas presented in this chapter allows for extracting useful information from a combination of almost any number and type of data sources. This information contains useful abstractions, that can then be used to make quantitative predictions. The data sources that can be thrown into the mix, and used for task-specific predictions, can be text (e.g. news articles, social media feeds, web search results), images, videos, psychometric measurements, econometric factors, stock market data, and many others. As long as the data can be represented in digital form on a computer, it can be used by deep learning systems in decision-making. AI, and in particular Deep Learning algorithms are prevalent in the background of a growing number of systems that make short and long-term decisions in finance. Additionally, some of the most prominent Deep Learning AI researchers are now directing the research labs at Facebook, Google, and Amazon. Since those companies house most of our data, the neural AI algorithms designed by those research labs will shape how our data is harnessed for various purposes in the future. In this chapter we introduce the general audience of alternative finance to some of the fundamental techniques used to extract useful information from such data. The appendix combines the ideas introduced within this chapter into a case study of an actual neural AI system design. Many of the ideas discussed here are already applied across a wide range of applications related to alternative finance. For instance, distributed representations of data such as vector space embeddings of words can be used to cluster people into groups based on their social media output. Such representations can be trained to be informative toward estimating their credit risk. Autoencoders can be used to derive informative features of noisy, high-dimensional data describing individuals. Such features can be trained to correlate with psychometric measurements and psychological theory. For example, in the context of group lending, the derived features can be made to form abstractions of individual credit default probabilities and psychologically informed parameters such as group cohesion, mutual trust,

cultural strength, and moral values. Hence, we can make use of psychological knowledge in a semi-autonomous framework, which tackles problems of higher complexity, which otherwise would have to be cut into sub-elements and laboriously studied by experts. Furthermore, deep learning systems can often find features and higher-order abstractions of raw data input, that aren't obvious to domain experts, or go beyond currently established psychological theories. Another example of how topics covered in this chapter are applied in industry is algorithmic trading. Later in the chapter we will see how gated recurrent neural networks and the encoder-decoder model, can be utilized for stock market prediction from various forms of data (e.g. recent news articles in combination with real-time price variations).

Machine Learning technology works in the background of an increasing number of software systems powering multiple aspects of modern society. In recent years a revolution on an unprecedented scale swept over AI community, that can be compared in magnitude to the fundamental conceptual shift that took place in theoretical physics at the beginning of the twentieth century (such as relativity theory and quantum mechanics). In the traditional Good Old Fashioned AI (GOF AI) approach, machines were meticulously hand-coded by domain experts to exhibit intelligent behavior by applying a set of pre-determined rules. This approach was extremely labor-intensive, couldn't scale to large problems, and wasn't robust to noise and new situations that haven't been considered by the engineers. After the long period of GOF AI, which centered on knowledge-based systems, with hierarchies of rules hard-coded by human experts, AI transitioned into the era of data-driven intelligence with the dawn of Statistical Machine Learning (SML). In this framework, the domain experts modeled the problem by formalizing its assumptions as statistical models with sets of tunable parameters. They also designed informative features of data that would be fed into the AI system. Learning then was understood as an iterative process of parameter adjustments, aiming at optimizing a chosen statistical criterion, which improved the system's performance at a given task. This approach was more robust to noisy data due to its probabilistic design. Furthermore it wasn't limited to a finite set of fixed rules of behavior, but instead could discover behavioral patterns from data. These systems enjoyed great success in many areas reviving popular interest in AI technologies. However, SML was still limited by the ability of human experts to design useful features of raw data, statistical priors, and parametrization of the model based on human intuition and knowledge. Implementing a successful SML system for real-world applications still required considerable human engineering and expertise. This form of AI was basically a version of applied statistics in disguise. SML-based AI

systems mostly consisted of a statistical computational model together with an efficient set of algorithms for fitting model parameters to various forms of data. Those systems could not be successfully applied to raw form data without significant feature engineering done by human experts. Hence, SML approaches aren't sufficient by themselves for being able to reach the holy grail of AI—the goal of achieving Artificial General Intelligence (AGI). Current decade saw a shift of AI to the domain of Neural Information Processing Systems. Artificial Neural Networks research started in the domain of computational neuroscience. Architectural improvements, together with learning algorithms tailored to neural systems, as well as advances in big data and hardware led to the current success of Deep Learning (DL), which refers to Machine Learning in the context of deep (multi-layered) neural networks with increasingly complex architectures. Deep Learning systems have beaten all previous approaches to AI, achieving State-Of-The-Art (SOTA) results across a diverse range of domains. The big success of Deep Learning systems comes from their contribution to representation learning—the ability of AI systems to be fed with raw unprocessed data, and automatically derive useful hierarchies of abstractions further utilized during learning tasks. With recent advances, neural systems are able to augment their own architectures and even learning processes based on data, further removing the need for human expertise. The key aspect of DL systems is that all of them can be trained regardless of task and type of data, using a unified general purpose learning procedure in an end-to-end fashion, with meaningful features being extracted automatically from raw-form data. Deep Learning is especially useful in situations with large amounts of unstructured, high-dimensional data present in business settings. Because of that, it is widely believed that novel applications of DL-based AI in settings outside of mainstream computer science will bring an avalanche of new developments across a wide range of applications relevant to alternative finance. For instance, fraud detection systems can use neural feature extraction to detect behavior that is hard to define explicitly.

Artificial Neural Networks (ANN) are computational architectures inspired by models of human brain. They consist of a graph of interconnected units, called neurons. Each neuron takes inputs from a collection of other neurons in the graph, and performs some operation, then propagates output (called activation), which becomes input to other neurons, or even back to itself (in case of recurrent networks). Neural AI models are trained on data using algorithms inspired by how functions are optimized in calculus (usually we want to minimize some definition of error). The core principle behind these learning algorithms is built on the idea of gradient descent—using the derivatives of the error function to iteratively approximate its minimum

value. However, instead of thinking of the error as a function of the data, we instead think of it as a function of the model (artificial brain) structure as defined by communication between sets of neurons. Hence, while optimizing the error function using gradient descent, we adjust the connections within the neural network—in other words, we alter our AI’s artificial brain structure to make it smarter (perform better) with respect to a particular task.

Research into artificial neural networks began in the early 1940s. The first model, called threshold logic, was developed by computational neuroscientists Warren McCulloch and Walter Pitts (1943). The first research results concerning neural learning mechanisms were published by a psychologist Donald Hebb (1949). Early ideas here go back to the beginning of nineteenth century, since primitive supervised neural nets were essentially variants on the linear regression theory of Gauss (1809, 1821) and Legendre (1805). Many design choices in artificial neural networks were inspired by neuroscience research. Early neural networks were based on two types of cells found in cat’s visual cortex (Hubel and Wiesel 1959, 1962). The first type of neuron, called *simple cell*, inspired sigmoid neurons. The second type, called *complex cell*, can be found in the primary visual cortex (V1), the secondary visual cortex (V2), and the Brodmann area 19 (V3). Complex cells were an early inspiration for deep convolutional neural networks, which led to the deep learning revolution due to their enormous success in the field of computer vision. First single hidden layer perceptron networks were developed in the 1960s (Joseph 1961; Viglione 1970). Deep multilayer networks of perceptrons were developed in the same decade (Ivakhnenko 1968, 1971; Ivakhnenko and Lapa 1965, 1967). First architectures that truly deserve the attribute deep were developed in late 1970s and early 1980s. Initial approaches focused on image recognition and introduced ConvNets (Fukushima 1979, 1980). Those first approaches didn’t use supervised gradient descent training, but instead relied on hard-engineering and some unsupervised techniques. Optimization of neural networks through gradient descent (Hadamard 1908) has been discussed since 1960s (Bryson 1961; Bryson and Denham 1962). Those first implementations of gradient descent were inefficient, and didn’t handle sparsity, recurrent connections, or links between nonsuccessive layers. Somewhat surprisingly, a 1969 book (Minsky 1969) on limitations of simple linear perceptrons with a single layer had a profound influence on the academic community, and discouraged researchers from further studying neural networks. Modern gradient descent optimization ideas for neural networks reemerged in 1980s (LeCun 1985; Werbos 1982). In 1986 Hinton published an influential paper (Rumelhart et al. 1985), which popularized an efficient back-propagation algorithm for training of deep neural networks.

It also showed empirical results suggesting the emergence of useful internal representations in hidden layers of deep architectures. This was the beginning of a wave of discoveries that formed a prelude to the modern deep learning revolution of past years. Recent breakthroughs include Long Short-Term Memory recurrent neural networks, attention mechanisms, dynamic memory networks, neural Turing machines, and deep reinforcement learning. We will discuss some of these ideas in more detail within this chapter. In the field of Natural Language Processing, the classical formal logic approaches, as well as count-based and SML methods in NLP were practically all replaced by DL models within the past years. Neural NLP systems now hold state-of-the-art results on all major natural language problems. The enormous success of Deep Learning in NLP was enabled in large part due to the work on regularization by researchers such as Yoshua Bengio, and by the advent of gated recurrent networks such as Long-Short Term Memory Recurrent Neural Networks (Schmidhuber and Hochreiter 1997). Similarly, the field of Computer Vision has been dominated by advanced convolutional architectures.

Learning in neural networks generally refers to episodic adjustments in weights associated with the network's edges, leading to incremental improvements in various measures of performance, commonly referred to as loss functions. Early efforts to develop learning algorithms for neural networks were inspired by neuroscience research pioneered by Donald Hebb. The Hebbian Theory introduced in his 1949 book titled *The Organization of Behavior: A Neuro-psychological Theory* (Hebb 1949) tried to explain the basic mechanisms for synaptic plasticity (the adaptation of neurons in the brain during the learning process). It was later used as basis for learning rules in Hopfield networks (Hopfield 1982). Key advance in learning algorithms for neural networks was the introduction of back-propagation algorithm (Williams and Hinton 1986). This approach is used in modern neural network research, and all modern AI systems are trained using some form of backprop. On a high level, back-propagation is a type of dynamic programming algorithm used to approximate the gradient of the loss function in the network's parameter space. We then perform some form of stochastic gradient descent to update the weights of the network in batches. Key property of this algorithm is that its time complexity is linear in the number of layers of the neural network, which matters for deep neural networks in particular. It involves a forward pass to compute the error, and then a backward pass to approximate the gradient. In the forward pass, data flows through the neural network from input to output. In the backward pass, derivatives of the error (gradients) obtained by comparing the model's output to the desired

response are computed for all parameters (connection weights between sets of neurons) in the network. Those gradients are then used to adjust the connection weights within the network in order to reduce future error. More details on this algorithm are given later in this chapter.

Deep learning refers to a field in neural network research concerning properties of complex, multi-layered, modular neural network architectures. The emergence of modern complex neural information processing systems created a paradigm shift in AI referred to by many as the “Deep Learning Revolution.” Even though neural network research spans several decades of the field known as “connectionism,” the modern Deep Learning trends in AI are quite recent, and can be traced back to the NIPS conference meeting of December 2012. During this (now arguably most prestigious) AI research conference, Geoffrey Hinton from University of Toronto (now also director of AI research at Google) made an unprecedented breakthrough in the field of computer vision using a deep convolutional neural network architecture. His contribution brought a larger improvement in AI vision systems than the collective efforts of the field over the previous two decades combined. Convolutional architectures were introduced to the field earlier by Yann Lecun (now a director of AI research at Facebook and a professor at NYU). Even though they were developed independently in computer science, their functionality has a striking resemblance to computational neural mechanisms present in human and animal brains, in particular within the visual cortices. Modern NLP systems incorporate a variety of deep architectures, such as ConvNet, LSTM-RNN, GRU, DQN, char-RNN, Memory Networks, and many others. A major contribution of deep neural architectures has been their ability to analyze data by inducing hierarchical representations, based on automatically extracted features. Higher level features are derived from lower level ones by the neural network when the information flows through consecutive layers (often entire modules composed of other deep networks). This leads to multiple levels of abstraction which are used to build up high-level concepts. Deep learning algorithms transform input data into distributed representations, where consecutive factors correspond to progressively higher levels of abstraction. Techniques used in deep learning are often inspired by neuroscience research. One example is neural coding—a study of the relationship between neuronal responses in the brain to various stimuli (Olshausen et al. 1996). For instance, an image can be represented as a vector of pixel intensity values at the input layer, then gradually processed into a set of edges, regions of various shapes, and even topological properties. In classical machine learning approaches to such problems, e.g. a Viola-Jones type object detection framework (Viola and Jones 2001), features are hand-selected (Haar

feature set in this case), and the algorithm uses statistical theory to derive the best classifier. In contrast, a deep convolutional neural network would extract a hierarchical set of progressively more abstract representations of images automatically. A type of neural network known as autoencoder is often used for feature extraction alone (Hinton and Zemel 1994; Bengio 2009). It is trained to reproduce its own inputs as closely as possible. After it has been trained to a desired level, we can use the information stored in the hidden units as a compressed representation of the training data.

The cross-pollination of AI and the field of neuroscience, as well as psychometrics (and even to some extent medical neurology) is on the rise. Several renowned deep learning researchers began collaborating with neuroscientists in a two-fold effort. In one direction, the insights gained through AI research are being used to understand how biological brains learn and process information. In the other direction, inspiration from neuroscience allows AI researchers to explore new algorithms and architectural designs for AI systems. After all, human brain is a working example of an intelligent system and consciousness in particular. Some of the most prominent deep learning researchers now work together with neuroscientists in this new interdisciplinary area bridging aspects of human and machine intelligence Bengio et al. (2016) and Sacramento et al. (2017, 2018).

This chapter introduces some of the most established ideas forming the foundations on which all modern AI systems are based. The core unifying principle behind Deep Learning approach to AI is the departure from discrete symbolic representations. Instead knowledge and reasoning are naturally emerging, dynamic manifolds of high-dimensional activity. This gives deep learning systems more of an “organic” feel, like biological systems, in contrast to the rigid clockwork mechanism of previous approaches to AI. Deep learning is a fast growing field with new discoveries made daily. Because it is a recent area of discourse that has not been properly organized yet, and often lacks theoretical foundations, it is hard to give a single comprehensive reference for state-of-the-art knowledge. However, for the same reasons, it is a challenging and exciting topic to investigate. New applications of deep learning systems to fields such as finance are especially promising as neural networks are surprisingly good at extracting signal from complex, noisy, high-dimensional data. Having understood the computational techniques and theoretical ideas discussed below, the reader should develop a deeper understanding of some of the foundational technical aspects of modern AI, and anticipate its future development and applications to domains outside of mainstream computer science. In particular, deep learning systems will become increasingly influential in finance over the coming years.

With the foundational ideas introduced here, the reader should gain a deeper understanding of how neural information processing is being harnessed to transform unstructured noisy data such as social media posts, into quantitative outputs such as credit score prediction, probability of financial fraud activity, or buy/hold/sell decisions of a stock market index.

2 What Is Machine Learning?

The most famous theoretical exploration of learning systems is the Vapnik-Chervonenkis theory of learning from data (Vapnik 2013). In some sense, all of learning can be understood as optimization (gradient descent on some error function). Mitchell defines learning in the following way (Mitchell 1997):

A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .

We can imagine a wide variety of tasks that fit into this rather broad definition. Task (T) defines the purpose of learning. For instance, in the domain of algorithmic trading, our task might be future index price prediction. Performance (P) is our measure of accuracy (e.g. how far our predicted prices are from the ground truth). Experience (E) is the data our system is allowed to learn from (e.g. past history of the stock market behavior). There are three main forms of machine learning: supervised, unsupervised, and reinforcement learning. The first two have access to the entire training set (either with associated answers or just the inputs), and reinforcement learning concerns agents collecting experience from the outside world through exploration. All three forms appear by themselves or in combination within modern AI systems. For instance, in the domain of Natural Language Processing (NLP), machine translation models are trained in supervised setting with source and target examples available. Image feature extraction or syntactic category induction in natural languages can be done in unsupervised way. Reinforcement learning can be used in algorithmic trading, where the agent is a program performing actions of buy/hold/sell within the perceived environment of the stock market.

2.1 The Bias-Variance Trade-Off

The art of learning relies on a balance between learning from the available experience while being able to generalize to unseen data. This is commonly referred to as the bias/variance trade-off. Learning is fitting parameters of a model to data. For instance we might model the probability θ that a given word appears in the context of the preceding $n - 1$ words (i.e. a probabilistic n -gram language model). The expected squared deviation of our model's estimate $\hat{\theta}$ of this probability can be decomposed into two additive terms:

$$\mathbf{E}[(\hat{\theta} - \theta)^2] = \mathbf{Bias}(\hat{\theta})^2 + \mathbf{Var}(\hat{\theta})$$

The bias term measures the expected deviation from the true value θ , while the variance term measures the deviation of our estimate from the mean averaged over possible samplings of the data.

We can think of this in terms of underfitting and overfitting. When the model underfits (high bias), it is usually too simple to capture the variation in the data (high training error). On the other hand, when a model overfits (high variance) it is learning the noise associated with the particular data used and is unlikely to generalize well (high test error) (Fig. 1).

3 How Much Data Do We Need?

It is usually hard to tell how much data is necessary to successfully train a machine learning system. It is even harder in case of deep learning, because we don't have much statistical theory available. In general, the answer to this question depends on multiple factors:

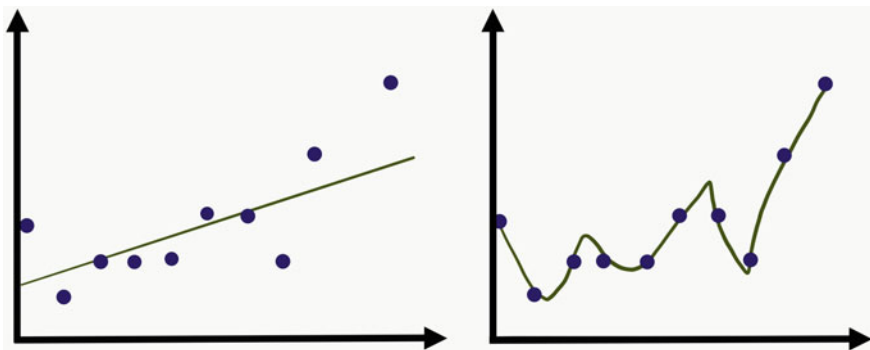


Fig. 1 Underfitting/high bias (left) vs. overfitting/high variance (right)

- complexity of the task (e.g. number of classes)
- number of model parameters
- number of input features
- quality of available training data (e.g. noise to signal ratio)
- properties of the loss function (how do we measure model performance)

It seems that on average deep learning requires more data, and continues improving in performance with the amount of data, while statistical models reach some level capacity earlier and stabilize, even if additional data is supplied. There is a saying in ML community that data is more important than the model itself. This means that simple models can often outperform sophisticated ones given large quantity of data to train on.

Learning high-dimensional problems like language is hard. In many cases the effective size of needed training data grows exponentially with the dimension of the problem. This is related to the “curse of dimensionality”—a term coined by Richard E. Bellman in the context of dynamic optimization. The basic idea is that most of the volume of a high-dimensional ball is concentrated near its surface, and this phenomenon is inflated with the dimension of the ambient space. Hence, in order to find informative parameter settings we need to search more of the range in each parameter.

3.1 Manifold Hypothesis

It is postulated that data generated by natural phenomena (such as human languages) comes from neighborhoods of lower dimensional manifolds within the ambient space (Fefferman et al. 2016; Narayanan and Niyogi 2009; Donoho and Grimes 2003; Belkin and Niyogi 2003; Brand 2003; Tenenbaum et al. 2000; Roweis and Saul 2000; Schölkopf et al. 1998; Cayton 2005). The resulting reduction in the number of effective parameters has a great effect on the efficacy of learning algorithms. If our data happens to be some noisy version of a lower dimensional sub-manifold, and our model architecture is naturally good at capturing that structure, we can mitigate the curse of dimensionality significantly. This is one of the possible reasons behind the unexpected success of deep learning. Modern AI systems based on neural networks can have millions of parameters. However, because of the regularities in the structure of naturally occurring data (in contrast to contrived problematic examples), they often converge with orders of magnitude smaller amount of data than theoretical considerations suggest. Similar phenomena might happen in the human brain. Perhaps, because neural information processing systems are loosely inspired by computational neuroscience

models of the brain, the design of an inter-connected network of artificial neurons is particularly good at capturing some properties of the naturally sourced data such as language or images. This makes sense if we consider the fact that our brain structure itself evolved in a lengthy optimization process to capture such data efficiently. Additionally, since our brain is designed to operate in this world, organic adaptation to naturally occurring patterns of data is a natural consequence of evolution, which can be understood as a form of long-term optimization. This suggests that biologically inspired algorithms might yield superior results when working with data containing patterns related to what human brain is tuned to process—e.g. areas that are connected with human behavior, in particular that of actors in the financial markets, or that of humans in social groups. These properties of deep learning AI make it a promising tool for alternative finance.

3.2 Synthetic Data Sets

There is a famous mantra in the machine learning community:

There is no data like more data!

We have seen that data scarcity is a real obstacle to successful implementation of modern NLP systems. If we can't obtain more data, a natural idea is to reuse the existing data. We can extend existing data by capturing some patterns, adding perturbation noise, and generate more samples that share some properties of the training set our model is trying to learn. Synthetic data has been used in machine learning research for some time (Bache and Lichman 2013; Graves et al. 2014; Weston et al. 2014; Mei et al. 2016), and recently in NLP (Weston et al. 2015) in the context of question answering.

Using synthetic data is also advantageous because it encourages Occam's razor approaches to research. There is evidence that given large quantities of data, researchers tend to develop simpler, more robust models (less bias), while data scarcity tends to imply complicated solutions tailored to the peculiarities of the training sets used (overfitting) (Halevy et al. 2009). Development of synthetic data often leads to deep insights into existing data, which leads to new algorithms, which then can be run on natural data with improved results.

Another form of synthetic data, especially in the context of reinforcement learning, comes from simulated environments. The case study at the end of

this chapter shows the process of training a designed AI model first on a simulator, then on more realistic data. This is a common practice in industrial development of deep learning systems.

4 Back-Propagation Algorithm

Majority of modern AI systems are now based on a data-driven deep learning approach. They learn to extract a hierarchy of features from training data by means of gradient descent optimization within neural networks. The algorithm used to achieve this is called back-propagation. We will review the basics of back-propagation learning below.

Suppose we have a neural network \mathcal{N} , which computes some function $f : X \rightarrow Y$. Given a set (called training set) $\mathbb{S} \subset X \times Y$ of correct input-output pairs, we would like to adjust the weights of \mathcal{N} in order to better match the training set. In order to define the meaning of “better” we first introduce a measure of network’s performance, called a loss function (sometimes the term error function is also used) $L^f : \mathcal{P}(X \times Y) \rightarrow \mathbb{R}$. In order for back-propagation to work, two basic assumptions about L have to be noted. First, it should decompose as $L^f(\mathbb{S}) = \frac{1}{n} \sum_x L_x^f$, where each L_x^f represents loss on a single training example $(x, y) \in \mathbb{S}$, and $n = |\mathbb{S}|$. Second, it should be a function of the outputs from \mathcal{N} . The loss function is used to compute how far the network’s predictions, for inputs in the training set, differ from the correct answers. An example of a loss function used in some settings is quadratic loss, which computes L2 distance between targets and predictions of the model

$$L^f(\mathbb{S}) = \frac{1}{2n} \sum_x (y - \hat{y})^2$$

where \hat{y} denotes the network’s prediction (i.e. $f(x)$) for the input example (x, y) . Other types of loss can be motivated by statistical or information theoretical considerations (e.g. we could compute KL-divergence between the distribution of the network’s outputs and the training examples). The back-propagation algorithm optimizes loss on the training set via stochastic gradient descent.

Here, we will derive the back-propagation update rules using a simple neural network model shown in Fig. 2. In the diagram of Fig. 2, squares denote input neurons (hence no activation function is applied there), and circles denote sigmoid neurons. Neurons are arranged in layers, which are denoted by named boxes I, J, K. We will use the following notation

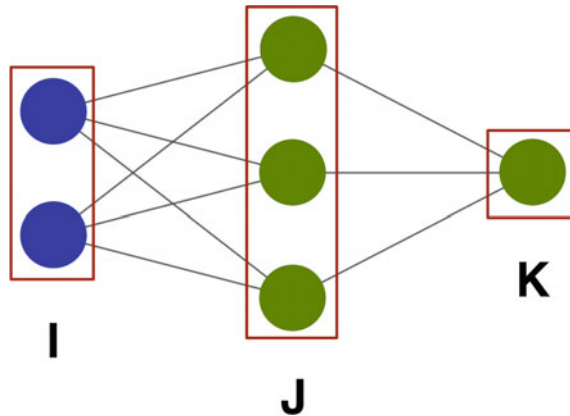


Fig. 2 A simple neural network

- x_j^l : input to node j of layer l
- W_{ij}^l : weight from layer $l - 1$ node i to layer l node j
- $\sigma(x) = \frac{1}{1+e^{-x}}$: sigmoid transfer function
- θ_j^l : bias of node j in layer l
- \mathcal{O}_j^l : output of node j in layer l
- τ_j : target value of node j in layer l
- $E = \frac{1}{2} \sum_{k \in K} (\mathcal{O}_k - \tau_k)^2$: loss

In order to derive the back-propagation algorithm, we need to compute the derivative of our loss with respect to weights and biases of each node in the network. We split our computation into two cases.

First let us consider the output layer (K).

$$\begin{aligned}
 \frac{\partial E}{\partial W_{jk}} &= \frac{\partial}{\partial W_{jk}} \frac{1}{2} \sum_{k' \in K} (\mathcal{O}_{k'} - \tau_{k'})^2 \\
 &= (\mathcal{O}_k - \tau_k) \frac{\partial}{\partial W_{jk}} \mathcal{O}_k \\
 &= (\mathcal{O}_k - \tau_k) \sigma(x_k) (1 - \sigma(x_k)) \frac{\partial}{\partial W_{jk}} x_k \\
 &= (\mathcal{O}_k - \tau_k) \mathcal{O}_k (1 - \mathcal{O}_k) \mathcal{O}_j \\
 &= \mathcal{O}_j \delta_k
 \end{aligned}$$

where we defined $\delta_k = (\mathcal{O}_k - \tau_k)\mathcal{O}_k(1 - \mathcal{O}_k)$ which is basically the difference between the example and the prediction, scaled by the derivative of the transfer function.

Similarly, we compute the gradient of the loss function for the hidden layer (J).

$$\begin{aligned}
 \frac{\partial E}{\partial W_{ij}} &= \frac{\partial}{\partial W_{ij}} \frac{1}{2} \sum_{k \in K} (\mathcal{O}_k - \tau_k)^2 \\
 &= \sum_{k \in K} (\mathcal{O}_k - \tau_k) \frac{\partial}{\partial W_{ij}} \mathcal{O}_k \\
 &= \sum_{k \in K} (\mathcal{O}_k - \tau_k) \sigma(x_k) (1 - \sigma(x_k)) \frac{\partial}{\partial W_{ij}} x_k \\
 &= \sum_{k \in K} (\mathcal{O}_k - \tau_k) \mathcal{O}_k (1 - \mathcal{O}_k) \frac{\partial x_k}{\partial \mathcal{O}_j} \frac{\partial \mathcal{O}_j}{\partial W_{ij}} \\
 &= \frac{\partial \mathcal{O}_j}{\partial W_{ij}} \sum_{k \in K} (\mathcal{O}_k - \tau_k) \mathcal{O}_k (1 - \mathcal{O}_k) W_{jk} \\
 &= \mathcal{O}_j (1 - \mathcal{O}_j) \frac{\partial x_j}{\partial W_{ij}} \sum_{k \in K} (\mathcal{O}_k - \tau_k) \mathcal{O}_k (1 - \mathcal{O}_k) W_{jk} \\
 &= \mathcal{O}_j (1 - \mathcal{O}_j) \mathcal{O}_i \sum_{k \in K} (\mathcal{O}_k - \tau_k) \mathcal{O}_k (1 - \mathcal{O}_k) W_{jk} \\
 &= \mathcal{O}_i \mathcal{O}_j (1 - \mathcal{O}_j) \sum_{k \in K} \delta_k W_{jk} \\
 &= \mathcal{O}_i \phi_j
 \end{aligned}$$

where we defined $\phi_j = \mathcal{O}_j(1 - \mathcal{O}_j) \sum_{k \in K} \delta_k W_{jk}$, which can be thought of as a weighted error at the output scaled by the derivative of the transfer function.

Because we can think of bias terms as connecting inputs always equal to 1, it is easy to verify that:

$$\frac{\partial E}{\partial \theta_l} = \begin{cases} \delta_l & \text{for } l \in K \\ \phi_l & \text{for } l \in J \end{cases}$$

The above derivation leads to the following algorithm for training the network from data:

- run the network forward on the inputs from the training set to compute its predictions
- for each output node compute $\delta_k = (\mathcal{O}_k - \tau_k)\mathcal{O}_k(1 - \mathcal{O}_k)$
- for each hidden node compute $\phi_j = \mathcal{O}_j(1 - \mathcal{O}_j) \sum_{k \in K} \delta_k W_{jk}$
- compute

$$\Delta_l = \begin{cases} \delta_l & \text{for } l \in K \\ \phi_l & \text{for } l \in J \end{cases}$$

- update:

$$W \leftarrow W_{ij} - \eta \mathcal{O}_i \Delta_j$$

$$\theta_l \leftarrow \theta_l - \eta \Delta_l$$

The above gives us the general rules for learning from data in the context of feed-forward artificial neural networks. In practice, we normally run the algorithm on smaller batches of examples instead of the entire training set, in what is called episodes. After repeating the training for a number of episodes we save the resulting weights, and then we can use the network to make predictions outside of the training set. There are various criteria we can use to determine when to stop. A popular choice is to stop when the change in loss between episodes slows down to some value. The learning rate, batch size, stopping criteria, are all examples of hyper-parameters. These aren't learned but are rather usually set by experts using some industry heuristics. There is a significant research effort aimed at automating hyper-parameter selection, but in general it is a hard problem, and good hyper-parameter choices are part of the art of deep learning.

We have derived a basic version of stochastic gradient descent for optimization of loss function in neural networks known as back-propagation learning algorithm. There are many modifications to this basic approach used in modern deep learning research. Some of them are inspired by physics to simulate momentum of a rolling ball on the energy landscape defined by the loss function, and require computing higher order derivatives of the error. We will not go into detail on those extensions here, but a good overview is given in Ruder (2016).

5 Vector Space Embeddings

Because inputs to neurons are numbers, deep learning systems need to preprocess various forms of data in order to encode it in numerical form before further processing can extract useful information. One could simply encode various data points with simple integer codes. However, this kind of representation isn't informative, since it loses reference to similarity between various inputs. Instead modern AI systems usually map data such as natural language text, sound, images, computer code, DNA, entire documents, or descriptions of human behavior and more structured inputs into a distributed representation inside a real vector space. This kind of representation has multiple advantages (such as compression and dimensionality reduction) but most importantly it comes with algebraic structure that can be exploited to perform meaningful operations on the data points. Different component dimensions in the ambient embedding space correspond to various features of the dataset. This section briefly describes the most basic methods for embedding sequence data with a shallow neural network, while the following section discusses deep recurrent embeddings able to learn long- and short-term dependencies in time series data prevalent in finance (e.g. performance of a stock index over extended periods of time). In case of natural language, the algebraic structure of the embedding space gives rise to linguistically meaningful operations (for example we can subtract semantic or syntactic properties of words by performing subtraction on their corresponding embedding vectors). In combination with more advanced neural architectures, this allows for various logical operations with natural language data. We can answer questions about the contents of input text, and perform transitional reasoning over facts contained within it.

One of the simplest and most successful modern NLP techniques is based on a feed-forward neural network with a single hidden layer. This method aims at embedding linguistic units into a geometric space with two main properties:

- Some measure of distance (e.g. L2 or cosine) should correspond to linguistic similarity (semantic, syntactic, etc.).
- Compositionality should have linguistic meaning. For instance, the path from “lecture” to “classroom” (i.e. the vector difference “classroom” – “lecture”) should correspond to a vector in a neighborhood of words such as “location”, and express the meaning of “where something occurs”. Hence, if we call that difference vector v , we should have that “reading” + v is close to “library”.

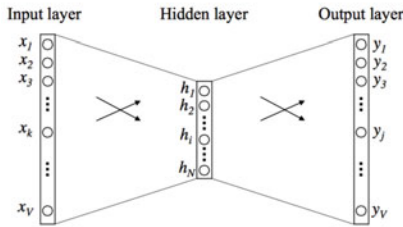
Multiple methods of computing distributed vector representations of words have been studied. The two most prominent examples are GloVe (Global Vectors for Word Representation) (Pennington et al. 2014) and Word2Vec (Mikolov et al. 2013a, b). GloVe is based on factorizations of the matrix of word cooccurrence statistics, and Word2Vec is based on a neural network encoding. In case of Word2Vec type of encoding, we have a single hidden layer feed-forward neural network with V input neurons and the same number of output neurons. The number V is set to the size of the lexicon. The hidden layer is meant to provide a lower dimensional explanation of the data, with $N \ll V$ neurons. There is a complete bipartite connection between successive layers (Fig. 3). The input is a sparse V -dimensional vector $x \in \mathbb{Z}_2^V$ with $\sum_i x_i = 1$ (i.e. there is a single 1 and all 0s in the input). The input vector corresponds to the words in the lexicon in some ordering, and the single 1 marks the word that is the current input to the network. The network in Fig. 3 represents a bigram model. It is straightforward to extend this to a general k -gram model by adding extra sets of V neurons to the input layer. We then train the network using back-propagation gradient descent algorithm on a raw corpus (e.g. Brown corpus). After training is complete, the hidden layer activations encode many syntactic and semantic properties of words in the corpus, based purely on their appearance in relation to other words. First, distance between vectors in the embedding space corresponds to syntactic and semantic similarity between the words they represent. This is because words with similar properties will appear in similar contexts, which means the words that can follow them are similar. Furthermore, the lower dimensional representation described by the network's weights allows for linguistically meaningful linear algebra operations. For example, in our experiments with a neural network trained using Google News corpus, the word corresponding to the vector that is nearest to the vector sum $\alpha(\text{'bigger'}) - \alpha(\text{'big'}) + \alpha(\text{'cold'})$ was found to be the word 'colder' (here α denotes the word embedding computed by the neural network). Many other examples such as $\text{king} - \text{male} + \text{female} = \text{queen}$ are famous in literature. Such effects are reminiscent of the psychometric tasks used to measure human intelligence capabilities. We can also observe interesting morphological phenomena. For example, taking several -ed forms of verbs, adding them together, and subtracting their null forms creates a vector that can generate new ed forms from null forms of verbs. Adding this vector to new verbs in null form tends to locate their corresponding past forms. There are likely subspaces of this ambient vector space correlating with various semantic, syntactic, and morphological properties of words.

Deriving word vector space embeddings from a large corpus allows us to produce a point cloud in a high-dimensional vector space representing words of the language. The Manifold Hypothesis in data science postulates that data generated by natural phenomena (such as human languages) comes from neighborhoods of lower dimensional manifolds within the ambient space (Fefferman et al. 2016; Narayanan and Niyogi 2009; Donoho and Grimes 2003; Belkin and Niyogi 2003; Brand 2003; Tenenbaum et al. 2000; Roweis and Saul 2000; Schölkopf et al. 1998; Cayton 2005). We can think of the evolving point cloud of words during training on increasingly larger amount of language data as noisy samples from the surface of some manifold (we'll call it the linguistic manifold) encoding the grammar of the language. Figure 4 shows side-by-side comparisons of a lower dimensional projection from vector space embeddings of French and English. The differences in shape result from the differences in grammar between these two languages.

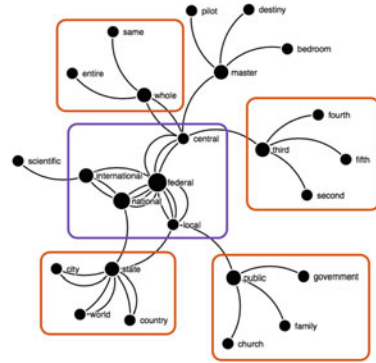
Word embeddings such as word2vec are often used for pre-processing natural language data at the input layer to more complex NLP systems. This is partially due to evidence in literature that they provide universal feature extractors for NLP tasks (Sharif Razavian et al. 2014). We often use it to represent the input as feature vectors at the input layer. Subsequent steps extract more task-specific features, and take sequence ordering into account (by means of other deep learning methods such as recurrence, attention, and convolution). It has been shown that vector embeddings of corporate quarterly reports can be used as a predictor for stock market variations in algorithmic trading systems. Another example relevant to alternative finance is analyzing language used by various groups by embedding large amounts of data into vector spaces and comparing their clustering properties. This allows us to analyze the communication of persons in the environment of the loan-taker, and infer the ingroup connection strength, which may act as a proxy for mutual support, and thus a predictive feature for credit default estimation.

6 Auto-Encoders

One of the goals of distributed representation, is dimensionality reduction. The original input might be a large discrete representation of words, and a data set of patterns ranging in billions of examples. The distributed representation compresses this input into a set of real vectors of fixed size (commonly around 300 dimensions), which encodes a large fraction of relevant information contained in the input. Such representations are often task-specific, which means they are trained to be most informative about achieving a



(a) neural network for computing vector embeddings (bigram model version)



(b) Nearest neighbors of words in a vector space embedding. Multiple edges mean that the nearest neighbor is already in the graph. Recursively expanding this graph produces clusters of words with similar semantic or syntactic properties.

Fig. 3 Word2vec

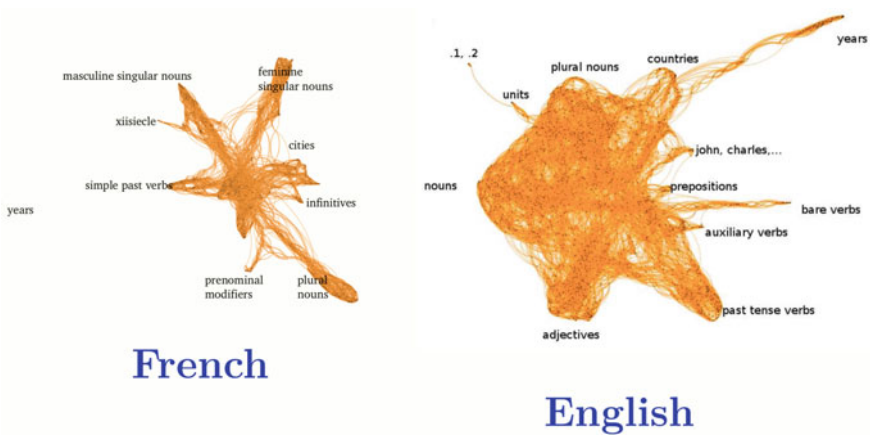


Fig. 4 Word vector space embeddings of entire languages (French and English). The shape of the resulting structure is related to the topological properties of the underlying linguistic manifold, which is influenced by the grammar of the language. We see that French and English grammars differ significantly by observing the differences in shapes of those projections

specific goal (e.g. predicting a sentiment of a product review provided by a user, or determining whether a quarterly report from a corporate executive board is positive or negative with respect to the future likely performance of

the business, and thus decide whether we should lend it money or buy its stock).

Another common way of achieving a lower dimensional representation of complex data is feeding it through an autoencoder neural network. In this case the obtained representation is task-agnostic, and is meant to produce a more general feature extractor that can later be used to preprocess inputs before feeding them to a variety of different models (which can be statistical machine learning models or more complex deep learning architectures). Autoencoding can be performed on a variety of data types (e.g. text, images, time series, or more structured information). The main idea is to embed the initially noisy, high-dimensional data into a lower dimensional vector space, which keeps enough signal that it can represent the data well for other purposes. This is usually achieved by a deep neural network where input and output layers are the same size, and intermediate layers extract abstractions of the input. The middle layers are usually much smaller size, and the network has a symmetric design (Fig. 5). The first half of the network is usually referred to as an encoder, because it encodes the initially large input into a compressed representation by propagating it forward through a lower dimensional “bottleneck” layer. This layer is usually thought of as “hidden” because it’s not directly visible when we imagine the autoencoder as a black-box that takes input and returns a reconstructed output. The second half of the network (to the right of the bottleneck in Fig. 5) is called the decoder, because it expands (decodes) the hidden representation of the input back to its original size. The network starts with randomly initialized weights, so the initial output is random noise. During the training process we perform unsupervised learning by feeding training pairs (x, x) to the network, where x is the data we wish to encode. The loss (error) function is designed to force the output of the network to be as close as possible to the input (therefore training examples are just pairs of the input, for each input in the dataset—i.e. the target of the output is the original input to the network). This process evolves the connection weight matrices between successive layers to encode the data into the hidden layer in such a way that this compressed representation can be used to reconstruct (decode) the original input with high accuracy. Intuitively, this means the hidden vector (represented by the hidden layer activations) needs to contain enough information so we can reconstruct the input with low error. Because this hidden layer is smaller than the original input, after enough training, we can use it to encode inputs into an efficient (in terms of size) embedding that’s informative of the main trends in the data. A trained autoencoder can then be used to encode new data of the same type (assuming it comes from a similar distribution) by feeding it through the

first half of the network and discarding the decoder part. The information encoded in the middle layer gives a lower dimensional embedding of new data points.

Suppose we have a high-dimensional noisy data such as psychometric measurements of employees at a company together with additional sources such as text describing their performance, numerical inputs representing their quarterly reviews by managers and a variety of other related inputs. We can use an autoencoder to derive a compressed representation of employees as fixed-dimensional vectors. Once we have this kind of representation, we can feed it to another AI model designed to classify employees based on their predicted performance in future assignments, or to match them to jobs or companies (which can also be encoded as vectors). Because we obtained a vector representation, we can also derive similarity measures to cluster the employees into groups for different purposes. Embedding discrete data into compressed distributed vector representations is a main component behind many emerging technologies applying AI to business management, especially in the HR sector. Alternative finance may benefit from such applications, since they allow for better internal decision-making. In the context of M&A this can be used for churn prevention and stabilizing the workforce. Modern acquisitions (especially in the technology sector) often focus more on talent than technologies, and deriving better methods to analyze workforce talent can be extremely beneficial in such settings.

In the following section we'll see how to deal with variable-size data (e.g. natural language text or speech, stock market prices changing in real time, time series data with long- and short-term temporal dependencies, etc.) in order to compress it into a fixed dimensional embedding.

7 Recurrent Neural Networks

Human brain deals with a constant sequence of sensory inputs. Because of this overload of incoming information and our limited memory storage, our cortices contain mechanisms to filter out relevant inputs and store them for later. Short-term memory in our brains can usually keep 5–9 time steps of information input. Relevant information is transferred to long-term memory in an augmented form, while less important information fades away. These neuroscientific facts inspired several AI architectures (among them are memory networks—which can learn to store relevant information and reuse it at a later time; and gated recurrent cells—which learn to encode long- and short-term dependencies over extended sequences of inputs). Here we

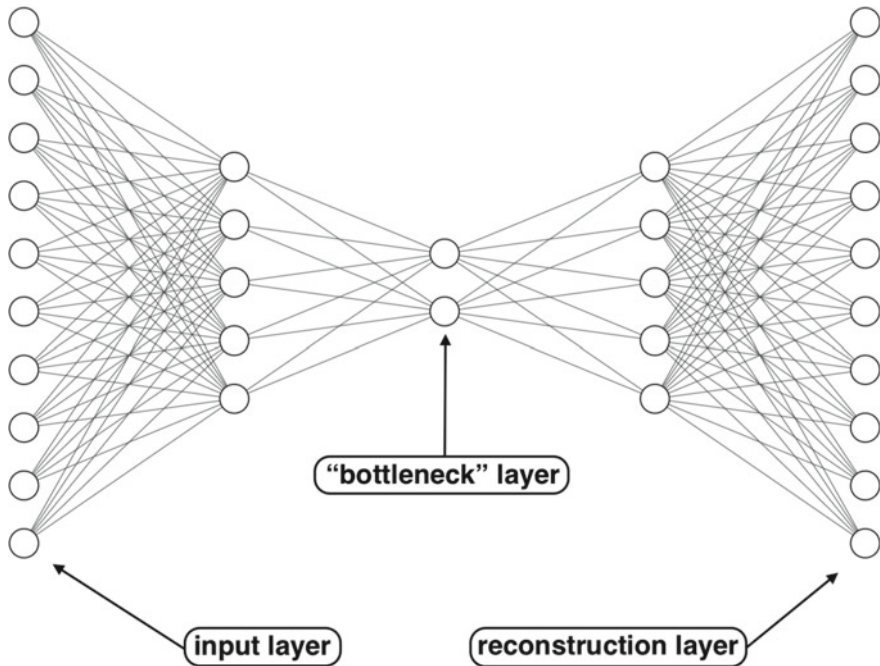


Fig. 5 Neural auto-encoder

focus on gated recurrent neural networks, which are particularly relevant to alternative finance applications. One of the most common data structures present in finance is time series data. For instance, variation in stock prices forms a stream of correlated numbers in temporal order containing long- and short-term patterns. The main method employed by neural AI systems dealing with time series data makes heavy use of the deep learning idea of recurrent connections and memory cells. Architectures based on simple feed-forward networks have no lasting memory of past inputs, other than the general knowledge that they extract from input-output pairs through periodical adjustment of their parameters. Recurrent neural networks address this problem. They allow information about previous inputs to persist, by feeding it back to themselves via recurrent connections. Because of this, they no longer form a simple acyclic graph, and their analysis is more challenging. We can conceptualize what's happening by unwinding the recurrent connections into an infinite graph formed of copies of our network (as in Fig. 6).

Remembering previous inputs is especially useful if they form coherent sequences such as strings of words in a sentence. Therefore, RNNs are particularly useful in natural language processing. In 1991 Jeffrey L. Elman from the departments of Cognitive Science and Linguistics at UCSD published a

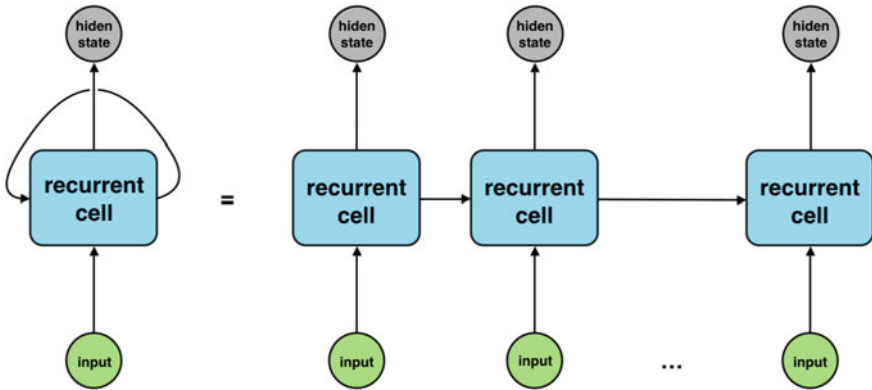


Fig. 6 Infinite graph representation of recurrent connection

seminal paper on distributed representations of text obtained with a simple recurrent neural network architecture. His paper showed that recurrent neural networks can learn long-term dependencies in sequences of words, and extract deep linguistic features of their input (Elman 1991). Elman's work was a first step in a sequence of important discoveries, most prominent of which include long short-term memory networks and gated feedback recurrent units (Schmidhuber and Hochreiter 1997; Chung et al. 2015).

In the context of text processing, a recurrent neural network takes as input a sequence of words in some vector representation w_t , where t is the index of a word in a sentence. Because words are fed into the network successively, t is often called the time factor. Let f denote the RNN cell, which can be a complicated function (such as an LSTM unit). In a single step at time t , an RNN cell takes the current input, and the previous hidden state h_{t-1} , and calculates the new hidden state h_t , which will be combined with the next input (e.g. consecutive word embedding in the text). Hence, for a single layer we have:

$$h_t = f(W \cdot h_{t-1} + V \cdot w_t)$$

for some learnable matrices W and V . In order to obtain a vector embedding of the entire input sequence, we just take the final hidden vector h_T . In case of multiple layers, in order to compute the hidden state in layer l at time t , we use the hidden state from the previous layer $l - 1$ at time t as input, and combine it with the current layer's hidden state from the previous time step:

$$h_{t,l} = f(W \cdot h_{t-1,l} + V \cdot h_{t,l-1})$$

For the initial input we use the word embeddings ($h_{t,0} = w_t$) as in the single layer case.

There is also a case of bi-directional RNN. In this model we read the input words in both directions, and obtain two sets of hidden states:

$$\begin{aligned}\vec{h}_t &= f(\vec{W} \cdot \vec{h}_{t-1} + \vec{V} \cdot w_t) \\ \overleftarrow{h}_t &= f(\overleftarrow{W} \cdot \overleftarrow{h}_{t-1} + \overleftarrow{V} \cdot w_t)\end{aligned}$$

We then use the concatenation of the two final directional state vectors, $[\overleftarrow{h}_1, \vec{h}_T]$, as the embedding.

Theoretically a plain version of RNN can solve problems with arbitrary temporal dependencies within inputs. However, in practice, especially if the gap between relevant information is large, it is very hard to train them to sufficient accuracy (Bengio et al. 1994). This is related to the unstable gradient problem, which is magnified in case of recurrent networks. Long Short-Term Memory recurrent neural networks (LSTM) (Schmidhuber and Hochreiter 1997) and Gated Recurrent Units (GRU) (Chung et al. 2015) (which is a special case of an LSTM) solve the unstable gradient problem in recurrent neural networks. LSTMs are composed of recurring cells that are able to retain important information and selectively forget some of the previous inputs. This is done through an internal circuit, with several types of gates. An input gate will add new relevant information into the state of the network, and a forget gate will filter out some information. The hidden cell state is passed on via recurrent connection and combined with input. This can be described formally by a set of equations (Fig. 7):

$$\begin{aligned}f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\ i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\ \tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\ C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t\end{aligned}$$

The cell state at previous time C_{t-1} is combined with current input x_t and previous output h_{t-1} , filtered via the input and forget gates, to form the new cell state C_t .

Since the original paper introducing the idea of long short term memory (Schmidhuber and Hochreiter 1997), several modifications have been proposed. We mention one popular version here. The gated recurrent unit (GRU) (Chung et al. 2014) combines input and output gates into a single “update gate,” and merges the cell and hidden state, among other

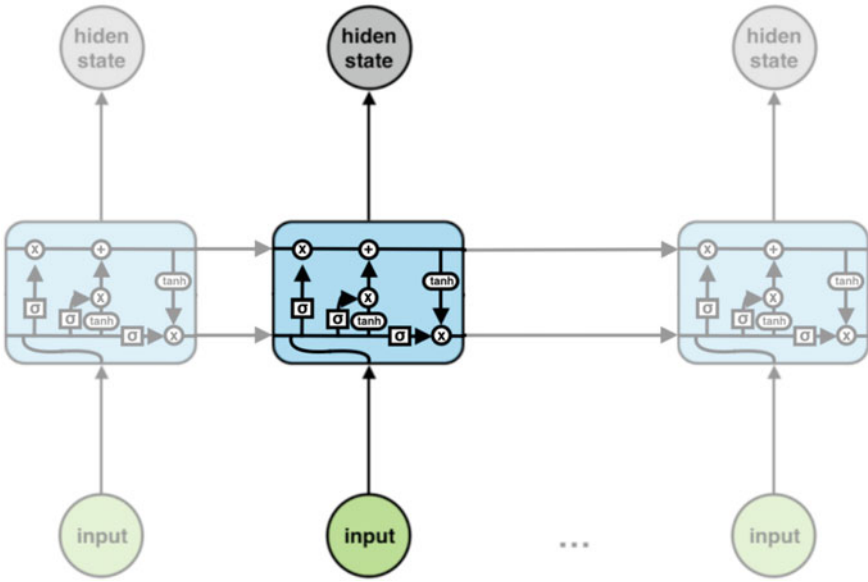


Fig. 7 Long short term memory network

augmentations (Fig. 8):

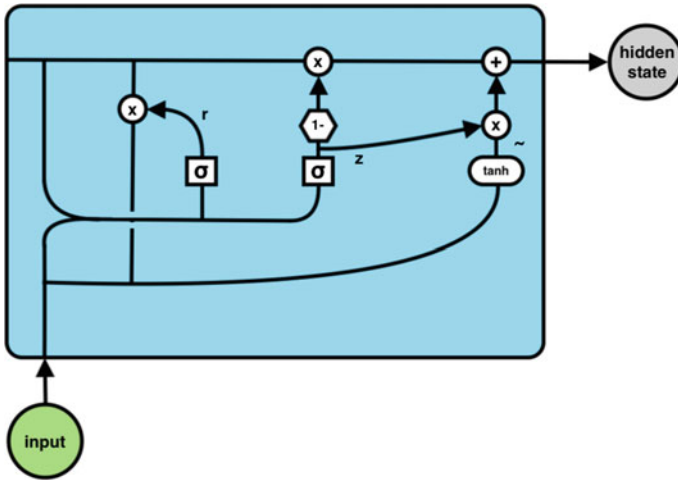


Fig. 8 Gated feedback recurrent unit

$$\begin{aligned}
z_t &= \sigma(W_z \cdot [h_{t-1}, x_t]) \\
r_t &= \sigma(W_r \cdot [h_{t-1}, x_t]) \\
\tilde{h}_t &= \tanh(W \cdot [r_t * h_{t-1}, x_t]) \\
h_t &= (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t
\end{aligned}$$

Extensive empirical comparisons of these and many more recurrent neural network architectures (see Greff et al. [2015] and Zaremba [2015]) reveal that they are all comparable in performance, and GRU is one of the best while maintaining relative simplicity and intuitiveness. Because of that, it has become increasingly popular since its publication.

8 Convolutional Architectures

One of the most important deep architectures is known as Convolutional Neural Networks (ConvNet). It got its name from a mathematical operation called *convolution*. One way to develop some intuition here is to think of convolutions in terms of probability density functions of independent events. Under this interpretation, given two PDSs f_A and f_B , describing distributions of independent random variables A and B , the PDF of their sum $C = A + B$, is given by the convolution of f and g .

$$f_C(x) = \int_{-\infty}^{\infty} f_A(y) f_B(x - y) dy = (f_A * f_B)(x)$$

Another good grounding for this concept comes from image processing. A common way to blur an image is to slide a localized function, called a *kernel*, through the image, and use it to compute the mean of the image under that kernel. We sum the values of the image pixels (first function) multiplied by values of the kernel (second function) at various positions, which is equivalent to convolving the image with the kernel. This method is common in statistics, e.g. non-parametric regression. In general a convolution of two functions f and g over domain (D) is given by:

$$(f * g)(x) = \int_{a+b=x|x \in D} f(a)g(b)da$$

or in discrete case:

$$(f * g)(x) = \sum_{a+b=x|x \in \mathcal{D}} f(a)g(b)$$

As we will see soon, we can think of convolutional neural networks as applying the discrete type of convolution to the input layer. First we need to introduce the CNN architecture. Normally, consecutive pairs of layers in a feed-forward neural network would form a complete bipartite graph. In applications such as image processing, networks are quite deep and consist of large numbers of neurons per layer. For instance, we would often have a separate input neuron for each pixel in the image. This leads to a very large number of distinct weight parameters we need to train from data. Convolutional architectures counter that by applying local filters at each layer via convolution. This is implemented by tying of weights in groups. So we only need to learn the tied groups of weights (features). For example, we could have a window of 9 input neurons (3×3 pixel square in the input image) connected to a single neuron in the first hidden layer (one per each group of four input neurons), with no other cross connections. So every neuron in the second layer would only communicate with (i.e. receive input from) four distinct neurons in the input layer. We end up with linear number of distinct weights, instead of polynomial in a standard feed-forward network. It also reduces learning complexity, because we only need to learn a smaller number of distinct connections (features), which are applied to different sets of inputs (e.g. different parts of an image). We can repeat the convolution multiple times. Hence, we again feed fixed groups of neurons from a hidden layer into the next layer, and repeat the process. This leads to hierarchical feature extraction, and has been shown to model increasing levels of abstraction (Bengio 2009). We often include several special types of convolutional layers, performing various statistical summaries of the features extracted so far. For example, a max pooling layer is very popular in computer vision. It simply fires if a certain feature is present in some region of a previous layer. This can be thought of as zooming out, because we get a coarser view of the inputs, which is more robust under local noise. Figure 9 shows a diagram of a simple 1D convolutional neural network with two neurons per window. Figure 10 shows a more realistic example of a 2D layer diagram with 4 neurons per window. Convolutional layers can get more complicated, and gates (like the pooling gate in our example) can even be entire neural networks with their own architectures (Lin et al. 2013).

Krizhevsky et al. (2012) introduced a wide array of new neural network techniques (including new convolutional architectures, regularization via

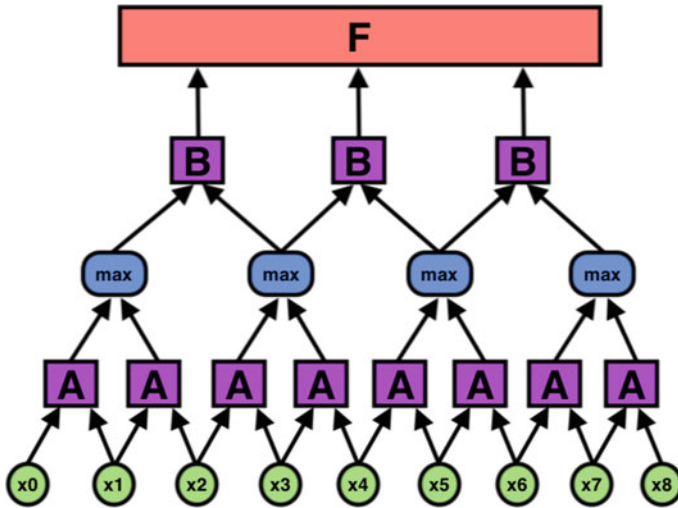


Fig. 9 Basic CNN diagram with max pooling

dropout, ReLU neurons, and the use of GPUs to train networks on large data sets) in their seminal paper on image classification. Their architecture (pictured in Fig. 11) is quite complex. Using this architecture, Hinton et al. were able to automatically extract a number of features for image recognition. Some of them are visualized in Fig. 12. The features learned include edges in various orientations (top of Fig. 12), as well as textures and color combinations (bottom of Fig. 12). Nobody designed those features by hand (as natural and useful as they appear for image processing). They emerged automatically as a result of back-propagation learning, from initially random weights. Furthermore, they resemble features extracted by the visual cortices in animal brains as convolutional neural networks were partially inspired by research in neuroscience. Representing images using such features instead of individual pixels reduces dimensionality of the problem, and makes the classifier robust to local noise (Fig. 13).

Convolutional neural networks have been applied successfully to problems such as neural machine translation (Gehring et al. 2017), search query retrieval (Shen et al. 2014), semantic parsing (Yih et al. 2014), sentence modeling (Kalchbrenner et al. 2014) as well as a variety of traditional NLP tasks (Collobert et al. 2011). Another interesting application is processing tensor data (such as images or higher dimensional matrices) in order to embed it into a real vector space of arbitrary dimensionality. This embedding can then be used to produce a probability distribution via softmax operation, or more interestingly as a starting hidden vector for another neural module

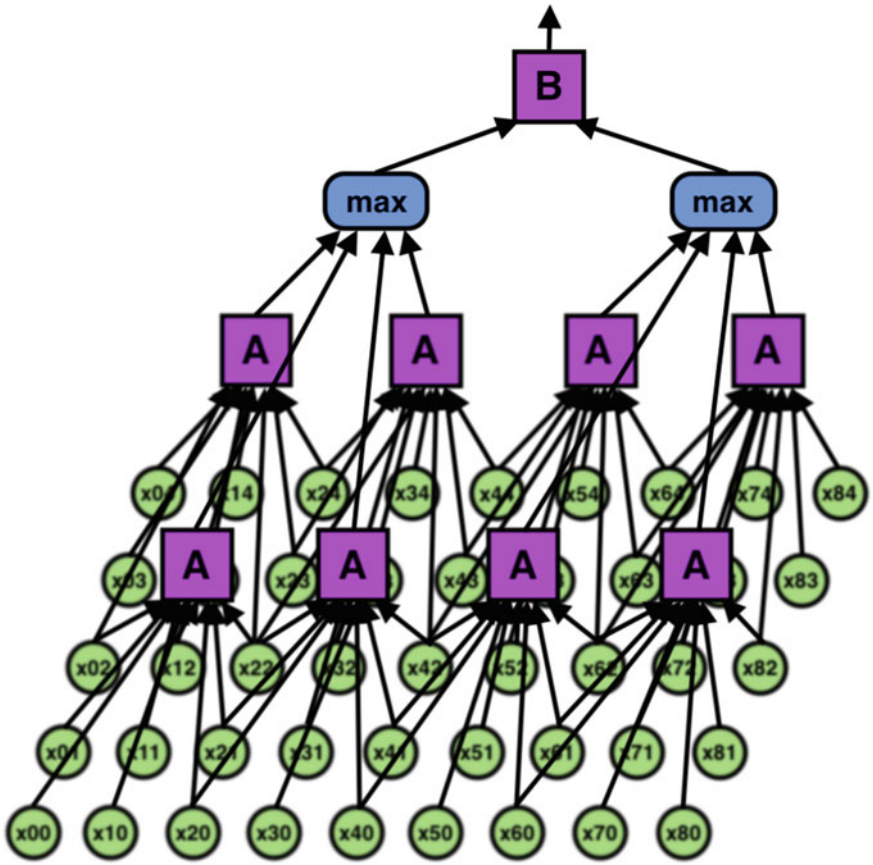


Fig. 10 CNN with max pooling for image processing

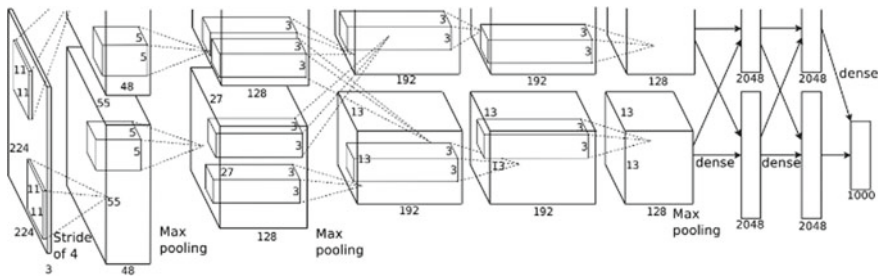


Fig. 11 Alex Krizhevsky, Ilya Sutskever, and Geoff Hinton—imagenet classification

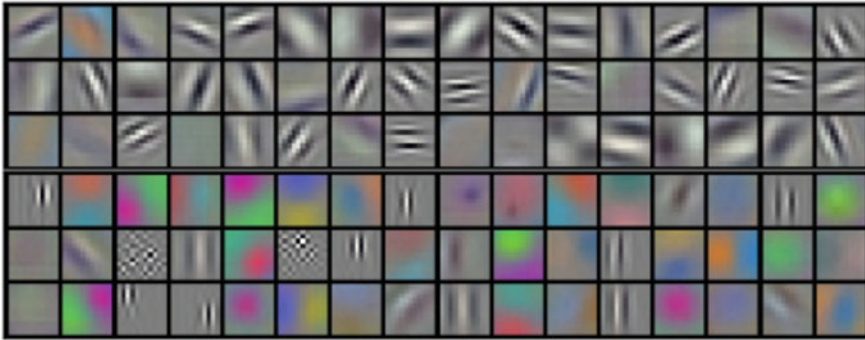


Fig. 12 Filters learned by the first convolutional layer on image classification task

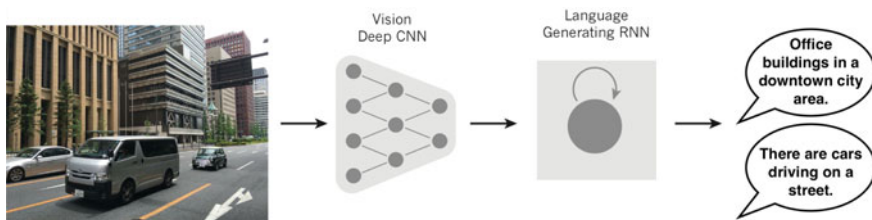


Fig. 13 Combining CNN and RNN allows relating various forms of data such as translating images into text

such as a recurrent neural network (instead of starting with randomly initialized weights). In this context, we can train such a hybrid system to associate images or other tensor data with temporal sequences such as strings of words. Famous applied example is image captioning, where this kind of system can write variable-length descriptions for still images as well as videos (Fig. 14).

9 Language Modeling and Transformer Based Architectures

The past two years (2018–2019) brought a surge of renewed interest in the field of Natural Language Processing due to several breakthroughs in deep neural language models. Those advances were based on ingenious uses of various forms of attention in neural networks. A seminal paper under the title “Attention Is All You Need” (Vaswani et al. 2017) introduced a new type of neural network architecture for extracting deep contextualized text representations from raw natural language data using a process based predominantly on repeated application of the “self-attention” operation in a



Fig. 14 Attention mechanisms can be used with both text and images to visualize the focus of the system on the pixel and character levels

model called the transformer. This kind of model transforms original vector space representation of linguistic units through a sequence of embedding spaces, where each successive mapping recomputes the representation of each token in context of its surrounding tokens. This allows for the semantics of words as seen by the neural AI systems to vary depending on the context and evolve over time. Such representations produce significant performance improvements on natural language understanding tasks. The AI research community was inspired by these advances and the field of NLP became the main focus of many prominent deep learning groups in industry and academia. This led to a series of publications in recent months that explored the capabilities of transformer architectures. The resulting language models became so good at generating naturally looking text conditioned on arbitrary prompts, that a major industrial AI research lab took a controversial decision to withhold the fully trained model weights, and the training code, from the scientific community on the basis that it would be too dangerous to release it to the public. This decision sparked a discussion about ethics and regulation of future AI research. In an unprecedented move, OpenAI—a company co-founded by Elon Musk focusing on AI safety—released a statement motivating their decision not to publish the SOTA model reported in their research paper:

Large, general language models could have significant societal impacts, and also have many near-term applications. [...] These findings, combined with earlier results on synthetic imagery, audio, and video, imply that technologies are reducing the cost of generating fake content and waging disinformation

campaigns. The public at large will need to become more skeptical of text they find online, just as the “deep fakes” phenomenon calls for more skepticism about images. Politicians may want to consider introducing penalties for the misuse of such systems, as some have proposed for deep fakes. Today, malicious actors—some of which are political in nature—have already begun to target the shared online commons, using things like “robotic tools, fake accounts and dedicated teams to troll individuals with hateful commentary or smears that make them afraid to speak, or difficult to be heard or believed”. We should consider how research into the generation of synthetic images, videos, audio, and text may further combine to unlock new as-yet-unanticipated capabilities for these actors, and should seek to create better technical and non-technical countermeasures. Furthermore, the underlying technical innovations inherent to these systems are core to fundamental artificial intelligence research, so it is not possible to control research in these domains without slowing down the progress of AI as a whole. [...] Nearly a year ago we wrote in the OpenAI Charter: “we expect that safety and security concerns will reduce our traditional publishing in the future, while increasing the importance of sharing safety, policy, and standards research,” and we see this current work as potentially representing the early beginnings of such concerns, which we expect may grow over time. This decision, as well as our discussion of it, is an experiment: while we are not sure that it is the right decision today, we believe that the AI community will eventually need to tackle the issue of publication norms in a thoughtful way in certain research areas. Other disciplines such as biotechnology and cybersecurity have long had active debates about responsible publication in cases with clear misuse potential, and we hope that our experiment will serve as a case study for more nuanced discussions of model and code release decisions in the AI community.

In an article associated with their published research paper on a then SOTA language model GPT2, they include “Policy Implications” and “Release Strategy” sections, where they make an argument for regulating AI research.

We also think governments should consider expanding or commencing initiatives to more systematically monitor the societal impact and diffusion of AI technologies, and to measure the progression in the capabilities of such systems. If pursued, these efforts could yield a better evidence base for decisions by AI labs and governments regarding publication decisions and AI policy more broadly. [...] we have had conversations with many external researchers, technology companies, and policymakers about our release strategy and the implications of increasingly large language models. We’ve also presented or discussed our work at events, including a dinner co-hosted with the Partnership on AI and a presentation to policymakers in Washington DC at the Global Engagement Center.

OpenAI along with several other research labs focusing on deep learning, are concerned that powerful NLP technology can be abused by corporate and state actors to spread misinformation and influence human behavior by producing high-quality fake news articles, comment section bots, product reviews, and even entire books at an unprecedented scale. On the other hand, some researchers are concerned that hiding such advanced technology from the public might be done under false pretenses to establish an oligarchy of powerful actors who have exclusive access to the technology. Another argument for releasing this technology to the public, is that researchers around the world will be able to test it, and find potential security risks and solutions. Otherwise the few who have exclusive access could weaponize it before the world is ready to defend itself. In his recent book, Max Tegmark, a prominent MIT professor, explores possible consequences of corporate research labs such as OpenAI or DeepMind (acquired by Google) developing advanced AI and keeping it secret. As a result of his research in this field, he co-founded the Future of Life Institute, whose goal is to inform regulators in the United States and abroad about the dangers of unconstrained AI development, and suggest policies to mitigate potential risks. The institute has been endorsed by multiple prominent business figures, politicians, and academics, including Elon Musk, President Obama, and the late Stephen Hawking. Several other advisory bodies on the intersection of industry, academia, and government have been founded around the world in recent months. Due to the enormous impact of AI technology on society, we see an emergence of a new interdisciplinary field involving scholars from social sciences, law, political science, economics, as well as computer science and engineering. This new field focuses on understanding the consequences of AI development on economic development, financial markets, labor market, political scene, and socio-economic aspects of future civilization. In Japan, several universities instantiated centers to deal with these problems and suggest solutions to the government, including Center for Cyber-civilization at Keio University in Tokyo. China's leadership under Xi Jinping, has formed similar initiatives, and Trump's administration held a summit on Artificial Intelligence for American Industry designating AI as a one of the top R&D priorities for the United States. Politicians around the world started paying attention to these issues, including Andrew Yang, a democratic presidential candidate, who put AI at the center of his campaign and is gaining traction in Silicon Valley (including an endorsement from Elon Musk).

In the following sections we introduce the basics of those new advances in NLP. This information should prepare the reader to follow recent developments and anticipate the near future. Regulators will benefit from gaining

insights into this emerging technology in order to make informed policy decisions.

9.1 Language Models

Language modeling is the most fundamental task in NLP—a field of AI dealing with natural language inputs. It is the basis of many solutions to a wide variety of problems involving AI systems with linguistic inputs. These tasks include (among many others):

- natural language understanding
- question answering
- machine translation
- document summarization
- dialog systems

The fundamental goal of language modeling is to assign high probability to utterances (usually sentences in plain text) that are likely to appear in data (i.e. belong to the language) and low probability to strings of words that are not. A trained language model can then be used to assign probability to arbitrary sequences of words. In the past language models were parametric statistical models estimated from data. However, they have since been replaced with much more successful deep neural network-based approaches. Generally, a neural language model is a neural network taking strings of words as input, and returning a scalar probability of those strings, which is trained to correspond to the likelihood that such a string conforms to a particular language, as induced from large quantities of text, often called a corpus:

$$\mathbb{P}(s) = \text{DNN}(\alpha(s))$$

where $s = w_1, w_2, \dots, w_n$ is a string of linguistic tokens (usually words), and α is some input embedding (usually into a distributed vector space representation via a gated recurrent encoder).

Normally, instead of thinking of a language model in terms of estimating joint probability of a string of words, we view it in terms of its ability to predict continuation of an input sequence. This is how language models are used in practice to generate naturally looking text. A seed state (random or an embedding of some desired information to condition on) is used to produce the first linguistic unit (e.g. a word). Subsequently, this output token is fed back into the model to produce the second token, and each successive token

is conditioned on the sequence of tokens produced so far. This can be naturally modeled with a gated recurrent architecture like an LSTM encoder. To obtain this interpretation, we apply the chain rule to decompose the joint probability of a string of words into conditional probabilities of generating a word following a sequence of words already generated:

$$\mathbb{P}(w_1, w_2, \dots, w_n) = \mathbb{P}(w_1)\mathbb{P}(w_2|w_1)\mathbb{P}(w_3|w_1, w_2) \dots \mathbb{P}(w_n|w_1, w_2, \dots, w_{n-1})$$

A neural language model therefore focuses on computing a conditional probability of word w_n following a sequence of words w_1, w_2, \dots, w_{n-1} . A simple model of this sort, used in many AI systems, computes this probability by feeding word vectors (e.g. from a precomputed word2vec embedding) into an LSTM encoder, which evolves the cell's hidden state vector. This vector can be used as a representation of the sequence. Another transformation (affine layer followed by a softmax or a more complex neural network) is then applied to this sentence embedding to produce the conditional probability distribution $\mathbb{P}(w_n|w_1, w_2, \dots, w_{n-1})$ over the tokens that can follow.

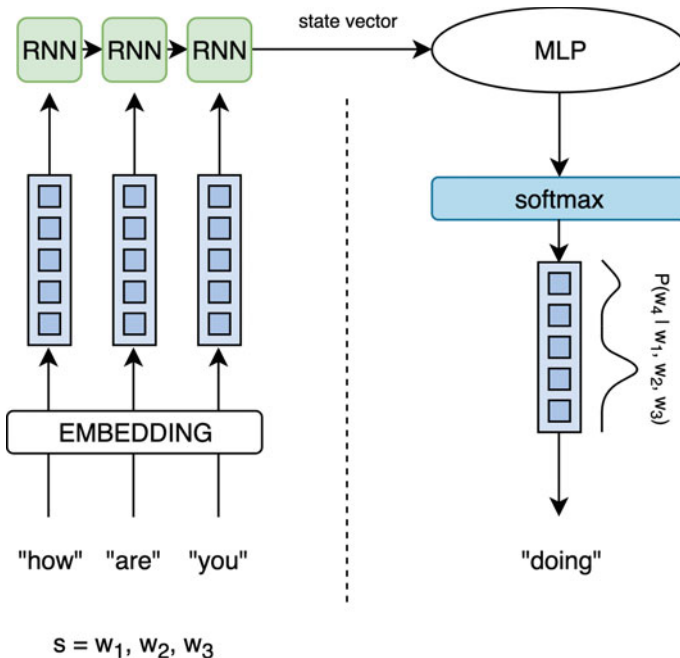


Fig. 15 Modeling a conditional probability of the next word in a sentence using a recurrent encoder

Textual output produced by recent deep neural language models can be very convincing. We can condition a decoder on an embedding of a human prompt, in order to start it off by generating text on a given topic. Here is an example from an OpenAI model published in 2019 (Fig. 15).

HUMAN PROMPT:

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

CONTINUATION GENERATED BY THE LANGUAGE MODEL:

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.

Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

Pérez and the others then ventured further into the valley. "By the time we reached the top of one peak, the water looked blue, with some crystals on top," said Pérez.

Pérez and his friends were astonished to see the unicorn herd. These creatures could be seen from the air without having to move too much to see them – they were so close they could touch their horns.

While examining these bizarre creatures the scientists discovered that the creatures also spoke some fairly regular English. Pérez stated, "We can see, for example, that they have a common 'language,' something like a dialect or dialectic."

Dr. Pérez believes that the unicorns may have originated in Argentina, where the animals were believed to be descendants of a lost race of people who lived there before the arrival of humans in those parts of South America.

While their origins are still unclear, some believe that perhaps the creatures were created when a human and a unicorn met each other in a time before human civilization. According to Pérez, "In South America, such incidents seem to be quite common."

However, Pérez also pointed out that it is likely that the only way of knowing for sure if unicorns are indeed the descendants of a lost alien race is through DNA. "But they seem to be able to communicate in English quite

well, which I believe is a sign of evolution, or at least a change in social organization,” said the scientist.

9.2 Transformers

The most successful recent approaches to language modeling are based on the idea of self-attention. It is similar to the attention in encoder-decoder models for sequence mapping, except it is done inside the encoder side on the input representations directly. The goal is to transform the initial input representation through a series of re-representation steps, where the embedding of each token is recomputed as a mixture of embeddings from its surrounding tokens. If the tokens are word vectors in a sentence, we can understand this as generating “contextualized” word representations. In this case, and in contrast to `word2vec`, the word vectors are not constant but evolve over a number of steps, where each word vector is re-expressed based on its context in the particular sentence being processed. Several issues arise when training deep transformer encoders, which include the necessity to encode positional information within the embedding space, and addition of “highway” connections that allow these networks to mix different levels of abstraction during the embedding process. The NLP/AI community has been very active on this topic, leading to multiple solutions emerging in the past months. At the time of writing this chapter, all top performing language models are deep transformer-based architectures. Detailed discussion of these methods is beyond the scope of this paper. Interested readers are encouraged to study the following methods:

- the transformer (Vaswani et al. 2017)
- embeddings from language models (ELMO) (Peters et al. 2018)
- bidirectional encoder representations from transformers (BERT) (Devlin et al. 2018)
- extended transformer through segmental recurrence (Dai et al. 2019)
- zero-shot task transfer in large language models (Radford et al. 2019) (Fig. 16)

10 Miscellaneous Topics

The three main architectural ideas (depth, recurrence, convolution) introduced in this chapter form a basis for many extensions, augmentations, and

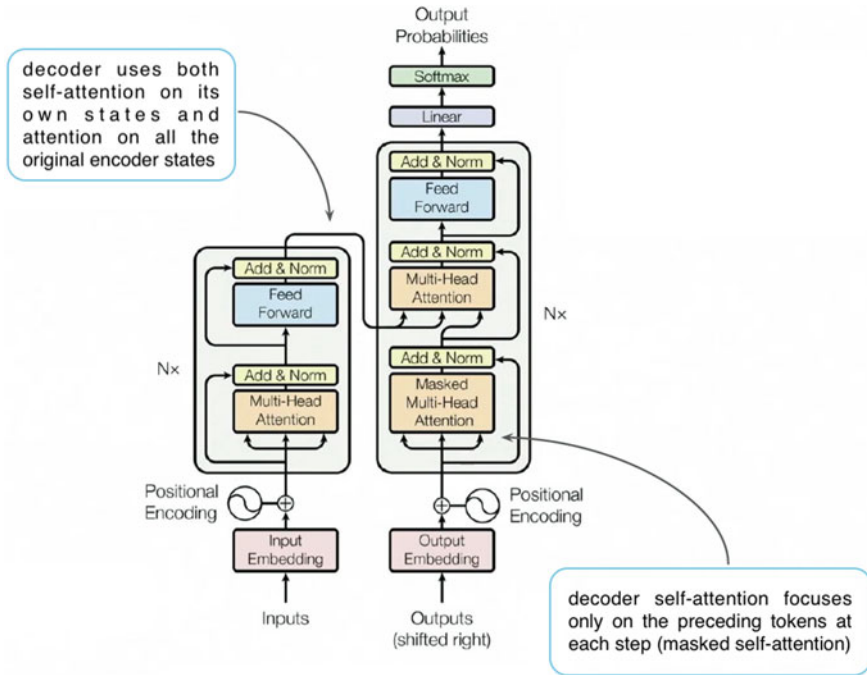


Fig. 16 Transformer architecture (single time step shown)

more advanced neural models emerging in modern AI systems. The methods introduced briefly in this chapter are the most established foundations of deep learning systems. However, in this rapidly evolving field, advances are made daily across the entire spectrum - from minor engineering tweaks (such as choice of an activation function or adding skip connections between layers), to deeper theoretical considerations on weights initialization or neural models inspired by ideas from statistics and mathematics of manifold theory (orthogonal weights, over-parametrization, importance sampling, conditional neural processes, generative adversarial networks, etc.). Several extensions of systems based on LSTM and CNN modules, such as attention mechanisms (which allow the network to focus with different intensity on relevant components of the inputs) or differentiable memory mechanisms (which allow networks to learn external memory management from data and act as a neural Turing Machine) can greatly increase the performance in practice. Architectures such as Dynamic Memory Networks (Xiong et al. 2016; Kumar et al. 2016) are inspired by ideas from psychology and neuroscience (functionality of the hippocampus, and the idea of episodic memory), and combine many of the architectures mentioned above into more complex modular systems.

A famous framework enjoying great success across a variety of applications ranging from vision to natural language processing is based on the idea of competition between submodules of a system. With Generative Adversarial Networks (GAN), a new class of algorithms was introduced (mainly by Goodfellow et al. [2014]) which revolutionized computer vision. At the core of this class of models are two competing neural networks (Fig. 17): the generator (G) and the discriminator (D). The generator attempts to generate naturally looking data, starting with random noise. The discriminator takes that generated “fake” data, as well as real data coming from some training set, and its job is to predict the probability that the given input is real versus one that’s generated by G. D tries to discriminate between those two possibilities by learning to assign a high score to real inputs and a low score to fake ones. Simultaneously G tries to improve at generating more realistically looking fake data in order to cheat D into misclassifying its inputs. This competitive dynamics is enabled by GAN’s loss measure, which rewards D for correctly distinguishing real examples from generated ones, while at the same time rewarding G for successfully confusing the discriminator. Because of this adversarial competition G learns the features of the real data faster and to a higher accuracy, than it would in a standard generative setting (non-adversarial). It can be shown that the Nash equilibrium of this kind of model produces a generator that learns the distribution of true data optimally, while the discriminator has to resort to random choice (guessing with probability of a half between data generated by G and true examples). This basic design forms the vanilla GAN model (i.e. the first version proposed by Goodfellow et al.). Many variants appeared in the computer vision literature since then, some of which are summarized in Fig. 18.

Another area not covered here in detail is the field of deep reinforcement learning, which was inspired partially by behavioral psychology, and approximates to some extent how humans learn by interacting with the world. In reinforcement learning the AI system is embodied in an agent interacting with (physical or simulated) environment. Instead of being given an explicit supervised training data, the agent learns by exploring the environment,

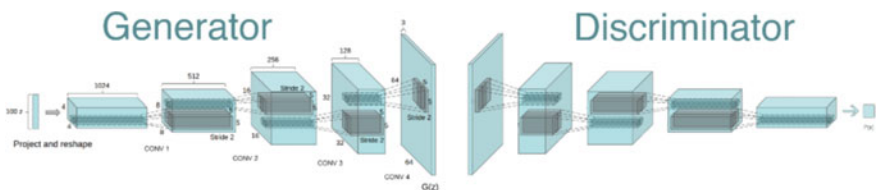


Fig. 17 Generative adversarial network

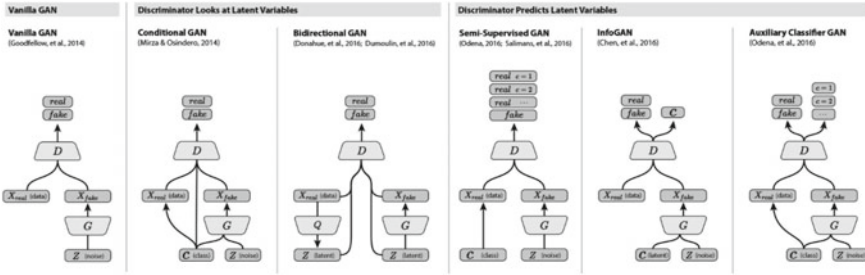


Fig. 18 Variants of GAN architectures

acting on it, and perceiving the effects of its actions. Through a mechanism of reward, which often comes at a later time after performing a sequence of actions, the agent generates its own semi-supervised experience, and adjusts its behavior to maximize expected rewards. Although reinforcement learning developed independently from deep learning in the past, recent attempts at merging these approaches to AI (e.g. by generating agents’ action policy via neural systems and using back-propagation to learn from experience with the goal of maximizing reward) lead to the hot new field known as “deep reinforcement learning”, which brought great successes in several areas making headlines around the world (e.g. AlphaGo—first AI software to attain super-human aptitude at the game of Go). The theory of reinforcement learning is beyond the scope of this chapter. Reinforcement learning is naturally suited to finance applications such as algorithmic trading, where a software agent could interact with the environment (i.e. the market) through actions such as buy/sell and collect rewards later in form of return on investment. Readers interested in this topic should consult the popular book on this subject by Sutton and Barto (2018).

Reinforcement learning combined with recurrent models producing sequences of symbols describing neural architectures in some formal design language lead to the idea of neural architecture search, where neural AI models can self-improve by producing recursive redesigns of themselves, where reward consists of increased performance of the mutated system on some task (Zoph and Le 2016). Such systems can also generate computer code (Neelakantan et al. 2015; Reed and De Freitas 2015) causing concern among some scientists and organizations about the future of AI, as future systems could evolve independently leading to superintelligence. Corporations such as IBM already use patented technologies evolved using similar methods, outperforming designs by human engineers. For instance, we might use such techniques to design circuits for physical devices. Several such

devices have been patented in recent years raising discussion among intellectual property law scholars. A famous example from software engineering is a redesign of the LSTM known as the NAS-cell. This redesigned RNN cell was developed automatically using Neural Architecture Search. The best performing cell generated using these techniques beats the LSTM type cells by a significant margin in multiple natural language processing tasks on benchmark datasets, and achieved state-of-the-art perplexity scores for language modeling on Penn Treebank at the time of its release. The PTB perplexity on test set for the NAS cell is 62.4, which is 3.6 perplexity better than previous state-of-the-art at that time. The NAS-cell also achieved new record results in character-level language modeling on PTB (perplexity score of 1.214). The resulting code was deployed by Google as part of TensorFlow, which is a mainstream tensor computation library used by AI industry to develop deep learning systems. This result is one among many in a growing trend focused on developing methods to automatically generate solutions to engineering problems. In algorithmic trading, we will likely see future trading systems designed by such automated processes. At some point this might lead to the stable markets, where the established players have equivalent trading strategies and nobody wins.

11 Final Remarks

This paraphrased quotation from Geoffrey Hinton—likely the most influential deep learning researcher in the world—summarizes the core difference between this new approach to machine learning in contrast to previous approaches to AI:

I think that people who assumed thoughts are symbolic expressions made a huge mistake. What comes in is a string of words, and what comes out is a string of words. Because of that strings of words were the obvious ways to represent things, and they thought that what goes on in between was some formal sequential language like a string of words. I think that what's between is nothing like a string of words. I believe that the idea that thoughts must be in some kind of language is as silly as the idea that understanding the layout of a spacial scene must be in pixels. However, what's in between isn't pixels or symbolic expressions. I think thoughts are these high-dimensional vectors that have causal powers - they cause other vectors, which is utterly unlike the standard view involving symbolic calculi.

Deep Learning approach to Artificial Intelligence suggests that thoughts aren't formal processes of symbolic manipulation, but instead are better viewed as evolving high-dimensional manifolds of neural activity. If we freeze time, the current mental processes taking place in the brain are fully determined by the neural communication throughout brain's synaptic pathways at that instance. Similarly, as a mathematical model inspired roughly by this structure, an Artificial Neural Network is defined by the strength of its synaptic connection represented by floating-point weights on edges between neurons forming the network. The set of all these weights (which can number in millions or higher) can be thought of as a point (or a point cloud) in a very high-dimensional real vector space (the neural network's parameter space). Under this interpretation, we can think of neural network states as points representing samples from some sub-manifold within this ambient space. During the learning process, the weights change as the network processes training data. This results in the vectors of neural activity moving around, and the resulting manifolds deforming until the learning process converges to some stable region.

Because of this manifold nature of neural networks, which are better understood as continuously changing dynamical systems, mathematical methods such as those employed by algebraic topology and differential geometry are likely to prove increasingly useful in future research on foundations of deep learning systems.

Due to the complex, noisy types of data, such as psychological factors of social groups, natural language contained in social media posts, social network connectivity graphs, tone of voice in speech, facial expressions captured by cameras during a loan interview, vital measurements of employees or their interaction during a workday (recorded by devices designed by companies such as Humanyze) and many others, separating signal from noise in order to extract what matters for a particular task, such as credit default prevention, is very hard. Deep learning offers an alternative approach in contrast to the theory-driven statistical models of the past. Via distributed vector embeddings of various inputs we can extract sub-manifold structure from big data. By applying various neural architectures, such as recurrent and convolutional neural modules, we can blend various types of data, and use it for decision-making in alternative finance. We can combine psychological theory with psychometric measurements of a group of people, as well as natural language data from their social media posts, in order to estimate their credit ratings, or decide if they are likely to be productive in a given job position.

In the appendix that follows, we discuss the design process of real-world neural AI systems that blend various forms of data, and are able to produce a variety of outputs ranging from probability estimates to arbitrary sequences of actions. The example discussed in the appendix grounds and crystallizes some of the core theoretical concepts discussed in this chapter. We also discuss how we can adjust the presented case study to other settings relevant to alternative finance.

12 Appendix: Implementing AI Systems

In this appendix we give some more concrete examples of how modern deep learning systems are designed and implemented in practice. These examples are by no means exhaustive of modern AI design principles, but rather serve as a grounding for reader's intuition, presenting the most common approaches. The selected examples make use of the basic concepts introduced in this chapter: distributed vector representations, convolutional neural networks, gated recurrent cells, and encoder-decoder architectures. After reading this appendix, the reader should get more grounded understanding of the theoretical concepts introduced in this chapter, and understand how to blend different techniques together to achieve a particular task of interest. For readers less comfortable with mathematical equations (which are given for completeness), we suggest skimming over them and focusing on the higher level overview. We include a rigorous case study resulting from applying the AI R&D process to a real-world problem.

12.1 Typical R&D Process for AI

Suppose you want to solve a particular problem of interest using deep neural networks, or perhaps automate parts of existing infrastructure. This task is usually performed by AI scientists, many of whom hold doctorates in computer science or related fields. Although the division of work between scientists who design the solution, and engineers who optimize it for eventual product release might differ between different organizations (e.g. academic labs, startups, government institutes, and large corporations), there are several core phases that are common across the spectrum.

In the planning phase, project managers (or principal investigators in research and academic settings) outline the problem to be solved, and some general characteristics of the desired solution. Those requirements often induce constraints on the approaches that the scientist will explore in looking

for a solution. In business settings they are often related to timeline and cost constraints set by the clients, as well as scalability of the planned system (which depends on the market). In academia or government research labs, these constraints might involve grant funding (which is managed by the PI), computational resources available (e.g. access to high performance machines), and time availability of the investigators (often postdocs or Ph.D. students who might be available for a given period of time). After the financial, computational, and time constraints are known, the scientists brainstorm possible approaches that would fit within those constraints.

In the brainstorming phase, scientists suggest various ideas based on their knowledge and experience. This is followed by a thorough literature review (which can take many weeks, even at smaller institutions), and a reading list is compiled. The purpose of the literature review is two-fold. First we want to avoid “reinventing the wheel.” Someone else might have worked on a similar problem in the past, and insights gained from their research might save significant time and money for the research team. Secondly, by following bibliographies of various publications on relevant topics, researchers often discover new ideas and suggestions that might be useful in solving the problem at hand. Promising publications are selected for in-depth study by the team. During the literature review process, the initially suggested ideas are ranked. Some are discarded completely, if for instance, publications by other scientists prove or suggest that such an approach is intractable. Once a couple of top ideas are distilled, the model design phase begins.

In the model design phase, known AI techniques are combined into diagrams representing the general architecture of the suggested approaches. We will see examples of these in the subsequent sections. These architectural choices at this stage are guided mostly by researchers’ intuition and experience in the field, as no empirical data is available yet. This intuition might sometimes come as inspiration from observations in other fields such as psychology and neuroscience, as well as mathematical considerations.

Next comes the implementation phase. This is elaborated in the following subsection. During the research process, scientists often choose a computation framework that abstracts away many technical details and allows for quick prototyping and testing of ideas. The most common language of choice is Python (which is commonly an interface to many efficient C libraries)—a multi-paradigm, high-level, general-purpose programming language popular in the AI community. Common choice of tensor computation frameworks includes PyTorch (developed by Facebook AI research team) and TensorFlow (from Google Brain team). TensorFlow is generally more appropriate for larger projects, or projects that need to be optimized for

mobile devices. PyTorch is increasingly popular in the NLP community for its intuitive design and flexible workflow. During the Implementation Phase, most promising ideas are tested on smaller dummy datasets, designed to test particular segments of the model.

After an idea has been successfully implemented and tested on simpler data, the Optimization Phase begins. Here improvements are made to the algorithms to make it scalable to real-world datasets, where efficient implementation is crucial to obtaining sufficient performance. This step sometimes involves systems engineers, who can rewrite research code in a lower level coding style.

After the final model code is optimized, the Evaluation Phase begins. At this stage, various experiments on real data are performed to isolate contributions of different components. Sometimes, design changes need to be made and the process described in previous steps repeats until satisfactory results are obtained.

The final phase depends on the final purpose of the project, and where it was developed. In academic setting the next phase would involve composing a journal article, that is likely published in a specific journal. In industrial settings, next steps would involve integration with existing applications or preparation of a production level deployment (which involves, among other things, adding more user-friendly interfaces).

In the following sections we will consider parts of this R&D pipeline exemplified with case studies. These examples are simplified to emphasize main ideas.

12.2 Implementation and Training—Tensor Computation Frameworks

A multitude of efficient tensor computation libraries is available in the Open Source community. Some of them were developed with corporate support from AI labs at Google (TensorFlow, Keras), Microsoft (MCT), Facebook (Torch), and Amazon (MXNet). Many others came from university research labs, government projects, and open source communities (DyNet, Chainer, Caffe and many others). These libraries are designed to automate many tedious, low-level tasks common to any deep learning project, such as data loading and pre-processing, saving and loading trained models, keeping track of and visualizing parameter changes, computing loss, and performing back-propagation of gradients through the network. Many frameworks come with higher level API for common design choices (such as commonly used arrangements of neurons, called layers). Deep learning libraries can be roughly

classified into two main types: dynamic (e.g. Dynet, Pytorch) or static (e.g. core part of TensorFlow). This refers to whether the computation graph of the neural network is determined before training, or whether the graph can be generated during runtime when training examples are provided to the model. Usually, the code is arranged in an OOP style, with a separate class for the model (which determines the network's architecture), data loader (which does the loading, pre-processing, and batching of training examples), and trainer (which accepts a model, data set, and hyper-parameters such as learning rate, then performs the tuning of trainable parameters on the provided data). In case of static frameworks like TensorFlow, there is also a separate object called a session, which actually performs the computation on the pre-determined graph. Let us briefly comment on the two most popular frameworks in the deep learning AI community.

TensorFlow is an open source software library for numerical computation using data flow graphs (Abadi et al. 2016). Historically, it is derived from an older framework called DistBelief, that was used at Google to implement their first deep learning models. It is especially well suited for machine learning applications, and it was developed at Google Brain for use in deep learning research projects. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. TensorFlow was originally designed as a static computation framework, meaning that the computation graph remains constant regardless of input, once it has been defined. This was a potential obstruction to developing some of the more advanced deep learning architectures, especially those involving recurrent neural networks in the area of natural language processing. However, later updates to TensorFlow introduced TensorFlow fold—an extension to the TensorFlow libraries that allows for training of neural networks with dynamic computation graphs. It is therefore now possible to define deep-learning models in TensorFlow that operate over data of varying size and structure, where the structure of the computation graph depends on the structure of the input data (Looks et al. 2017). TensorFlow is supported on a wide array of devices and architectures, including mobile devices, and cloud computing services. It also comes with an array of visualization tools grouped under TensorBoard visualization package. It is a rich visualization tool which can be used in debugging, or to gain inspiration during neural network research.

TensorFlow's (along with its high-level API called Keras) popularity in academic research grew exponentially fast since its release, and it is now the most frequently cloned repository on GitHub in the machine learning category (Fig. 19). Because of this strong community support, especially among

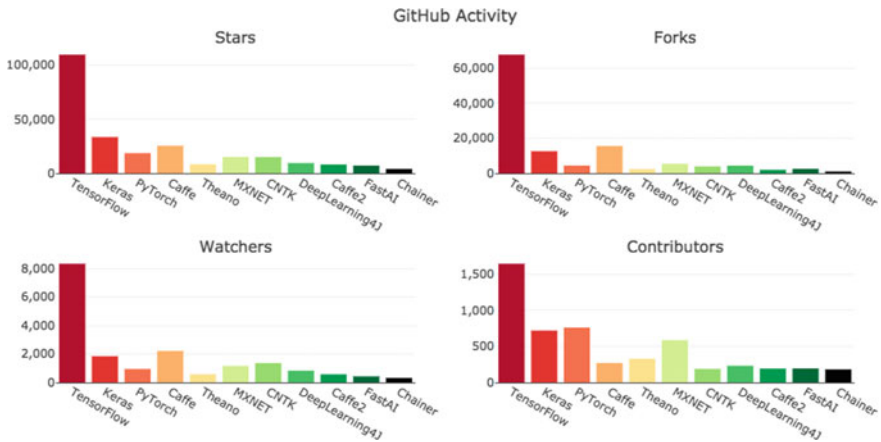


Fig. 19 GitHub activity statistics for top Deep Learning frameworks in 2018

academics, and corporate support from Google, it is easy to collaborate with people on projects, and there is already a large code base to recycle, which saves time and allows developers and scientists to focus on novel ideas instead of “re-engineering the wheel.”

TensorFlow has many merits but it also comes with significant downsides, especially when working on NLP applications. In recent years many deep learning researchers have switched to a newer framework called PyTorch, which is much better suited to work with language data. Two main issues with TensorFlow that PyTorch alleviates are static computation graph and reduced transparency when debugging complex code written in Python. PyTorch integrates smoothly with native Python style of writing programs, and the NumPy package for scientific computing (providing powerful n-dimensional array operations, sophisticated broadcasting rules, and a variety of useful mathematical operations). PyTorch has been an ongoing open source project led by scientists from Facebook AI Research. It is based on the ideas from Torch, a popular neural network library, rewritten and redesigned with Python in mind. It provides an efficient and dynamic tensor computation framework for research and prototyping of neural network architectures. Through the support of a dedicated community of developers, many of whom are deep learning researchers working in NLP, it is quickly becoming the top choice for work involving neural networks in the context of natural language. Pytorch was designed to improve on TensorFlow in the areas most relevant to researchers, allowing quick prototyping of ideas. The goals of TensorFlow are more aimed toward heavy industrial applications and

embedded systems, which need to be highly optimized. This ability to optimize code comes at a price of reduced transparency and cumbersome design patterns. Because of new open data standards for sharing neural models, it is possible to export PyTorch design to a different framework for generating more efficient code optimized for large-scale applications. However, the main concern of PyTorch providing a better workflow for research and developer efficiency.

Describing engineering details and coding practices for deep learning frameworks is beyond the scope of this chapter. However, the official project websites for TensorFlow (<http://tensorflow.org>) and PyTorch (<http://pytorch.org>) come with detailed documentation and common usage examples. I recommend the interested readers skim through the official tutorials available on the corresponding project sites.

12.3 Case Study: Combining Multiple Time-Series Data to Produce Temporal Sequences of Predictions

In this example, we examine an example of AI system design combining high-dimensional temporal sequences of inputs in the form of natural language text, as well as visual inputs in form of images, in order to produce a third sequence of outputs. This model can be easily augmented to work with stock market data (discussed later). However we show a case of natural language and visual inputs in order to show how the main ideas discussed earlier in this chapter can be applied in a real-world setting to combine different types of data. The techniques used here exemplify the main ideas from our introduction to neural information processing systems: distributed representations of data, dimensionality reduction, gated recurrent networks, the encoder-decoder model, attention mechanisms, and convolutional architectures.

In order to exhibit the interaction of those components in an applied scenario, we will study the problem of instruction following in the context of autonomous vehicle navigation (e.g. a self-driving car).

Natural language is the most intuitive and flexible way for humans to communicate. Unfortunately when issuing instructions to machines, humans are forced to master artificial, over-constrained, domain-specific ways of expression. Diverging from those specifications often renders instructions incomprehensible to machines, or worse, leads to unexpected responses. Because of this, significant effort has been dedicated within the AI community to develop systems bringing us closer to effective natural language communication between humans and machines. This fits into a bigger effort at multiple institutions, directed toward natural language understanding.

Natural language provides a rich, intuitive, and flexible medium for humans and robots to interact and share information. Apart from convenience, there are many situations in which being able to instruct machines in unconstrained natural language is a necessity.

Robots assisting first responders in search and rescue missions, need to be able to communicate with victims and personnel untrained in domain-specific languages. For example, a robot might engage in the following dialog when it encounters a victim:

- Robot: Someone is on the way to get you out of here. Are there any other people around who need help?
- Victim: I saw someone in the main lobby.
- Robot: Where is the main lobby?
- Person: Exit this room and turn right. Go down the hallway past the elevators. The lobby is straight ahead.
- Robot Understood.

Another example is instructing a self-driving car which route to take, or teaching it a location not present on its map.

- Passenger: The map is wrong, my house is actually on the other side of this building. Take left, then right, and my house will be next to the large tree.
- Car: Understood. Updating map location.



Fig. 20 Human-machine communication in natural language

In this case study we are interested in translating path descriptions given to a robot in natural language, into sequences of actions that robot can follow. This is an important practical problem for a variety of reasons. First, following directions requires the ability to understand spatial language. Because spatial language is pervasive, this ability is important for almost any application of natural language to robotics. Second, a system that understands directions is useful in many scenarios, including health care, companion robots, search and rescue, and self-driving vehicles carrying passengers. Third, it is natural to ask humans to create a set of directions through an environment, yielding an open-ended yet task-constrained corpus of language. Finally, there is a natural correctness metric when evaluating a robot's performance at following natural language directions: did it reach the correct final destination?

The goal is to relate a sequence of words representing a natural language command, to a sequence of actions for a robot controller to execute. Grounding of natural language instructions for human-machine interaction, such as directions issued by passengers of self-driving cars, has mostly been approached using graphical models and statistical machine learning approaches with hand-crafted features. The main obstruction to the application of deep neural network models for this purpose is scarcity of quality training data. In order to approach this problem in a data-driven manner, we generate a large data set from crawling Google Street View over randomly generated paths in a city. We obtain location and image data from Google Maps, and use image classification and language generation techniques to derive a corpus of natural language instructions paired with environment observations and robot action sequences. This data is used to train an alignment-based encoder-decoder model with Long Short-Term Memory recurrent neural networks (LSTM-RNN), which translates natural language instructions to action sequences based upon a representation of the observable world state (Fig. 20).

Research efforts on algorithms allowing free-form natural language communication with machines (MacMahon et al. 2006; Kollar et al. 2010; Chen 2012; Chen and Mooney 2011; Kim and Mooney 2012, 2013; Kong et al. 2014; Hemachandra et al. 2015) have mostly focused on the idea of symbol grounding defined in Harnad in the early 1990s (Harnad 1990), which concerns associating linguistic terms with the objects (physical or abstract) that they describe. Early solutions to the symbol grounding problem used hand-crafted mappings (Tellex et al. 2011; MacMahon et al. 2006). Statistical methods have also been applied to infer the meaning of words in the context of perceptual inputs (Mooney 2008; Matuszek et al. 2010). These methods require human engineered features and annotated corpora. Another

direction pursued previously is to treat instruction following problem as that of learning a parser defining a mapping of natural language into its formal equivalent, that can be processed by the machine. Chen and Mooney (2011) use probabilistic context free grammar (PCFG) induction to learn groundings for a learned lexicon, while Artzi and Zettlemoyer (2013) and Artzi et al. (2014) use a combinatory categorical grammar (CCG) based semantic parser to map natural language into a lambda calculus representation. Simultaneously, a parallel set of approaches have been developed which can be grouped under a label of probabilistic world model grounding. These techniques map natural language into sets of corresponding objects, locations, and actions within the agent's world representation under a probabilistic model of symbol to world correspondence. Interpreting instructions in this framework is performing inference in that learned model. Several previous approaches aimed at restricting the complexity of the search space by adopting a probabilistic approach with simplifying assumptions of independence. Kollar et al. (2010) developed a generative model of spatial relations, adverbs, and verbs. In particular graphical models (factor graphs) were used, which represent a factorization of the conditional probability distribution of groundings, given the sequences of words representing commands given by a human. Tellex et al. present a discriminative model that captures hierarchical and compositional structure of language. Such factorizations can be derived automatically by parsing natural language commands into Spatial Description Clauses, and then using the structure SDC decomposition of a command to induce the corresponding factor graph, with random variables representing words, groundings, and correspondence variables (which describe if groundings are correct).

Instead we harness a new array of powerful techniques for working with sequence data that have emerged in the neural network community. Deep learning approaches to sequence-to-sequence mapping have been increasingly successful with application in machine translation (Sutskever et al. 2014; Bahdanau et al. 2014; Cho et al. 2014), natural language generation (Rush et al. 2015; Wen et al. 2015), and image captioning (Kiros et al. 2014; Mao et al. 2014; Donahue et al. 2015; Vinyals et al. 2015; Chen and Lawrence Zitnick 2015; Karpathy and Fei-Fei 2015). We propose a deep learning approach to this problem, which does not require contrived formal languages such as SDC. In contrast to those prior methods, the deep learning model uses no specialized linguistic resources (e.g., parsers) or task-specific annotations (e.g., seed lexicons). It is therefore generalizable to a variety of human-machine interaction settings. The mapping is learned automatically

in an end-to-end fashion, purely via the means of back-propagation in a deep recurrent neural network.

In order to allow machines to interpret natural language instructions, we need to map sequence of words in a human language representing the instructions, to the appropriate sequence of actions corresponding to commands that robot's actuators can execute. On the high level of abstraction, it is therefore an instance of a sequence-to-sequence learning. However, the robot also needs to disambiguate the instructions in the context of the world environment it operates in. Therefore, we also need to consider a third sequence—the temporal sequence of perceptual inputs received by the robot in the process of executing the instructions.

To develop a solution to this problem, we approach instruction following as a form of neural machine translation (Bahdanau et al. 2014). In this setting instead of translating, for example, from English to Japanese, we translate from English to “machine language,” i.e. to the sequence of instructions that the machine can execute. However, in contrast with simple machine translation, we also consider the percept sequence, and use it to alter the interpretation of the input instructions in real time. First, the robot accepts the human instruction, and develops a memory vector representing its general understanding (encoder step). Afterward, robot takes an action based on that initial understanding, and receives new input from the environment (e.g. if the robot moved, a new image from the cameras will be processed). At each time step, another action is taken and a new world observation is received. This iterative process augments the initial memory vector representing the natural language instruction originally give to the robot. Eventually, a stop action is emitted. We can then evaluate how accurately the given instructions were processed. For instance, if we consider directions given to a self-driving car, we can measure the distance of the final location of the car from the desired destination.

Abstractly, the goal is to derive a model over sequences of actions conditioned on world state and natural language instructions: $P(a_{1:T}|y_{1:T}, x_{1:N})$. Then extract the maximum probability sequence of actions under the trained model.

$$\begin{aligned} a_{1:T}^* &= \operatorname{argmax}_{a_{1:T}} P(a_{1:T}|y_{1:T}, x_{1:N}) \\ &= \operatorname{argmax}_{a_{1:T}} \prod_{t=1}^T P(a_t|a_{1:t-1}, y_t, x_{1:N}) \end{aligned}$$

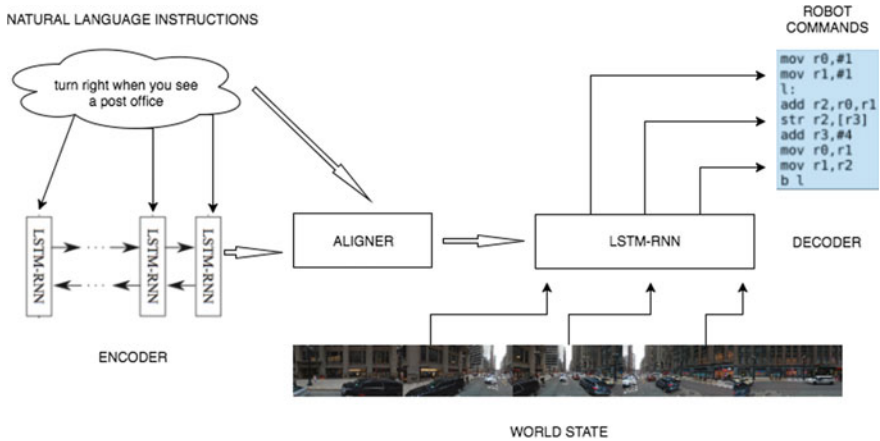


Fig. 21 Instruction following architecture

We approximate this conditional probability using a neural encoder-decoder architecture. During the encoder stage, we use a bidirectional recurrent Long-Short Term Memory neural network similar to that of Graves et al. (2013) to derive a sequence of hidden annotations $h_{1:N} = (h_1, h_2, \dots, h_N)$, where we think of annotation h_i as a summary of the first i words of the natural language instruction given to the robot.

$$\begin{pmatrix} i_j^e \\ f_j^e \\ o_j^e \\ g_j^e \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} T^e \begin{pmatrix} x_j \\ h_{j-1} \end{pmatrix}$$

$$c_j^e = f_j^e \odot c_{j-1}^e + i_j^e \odot g_j^e$$

$$h_j = o_j^e \odot \tanh(c_j^e)$$

Here e designates that the variables correspond to the encoder phase, σ is the logistic sigmoid function, T is an affine transform, i, f, o are the input, forget, and output gates of the bi-LSTM, and c is the cell state activation vector. The cell memory is updated iteratively based on previous memory and current input under regularization induced by the input and forget gates. The use of a bidirectional LSTM-RNN as the encoder is inspired by its success in speech recognition and machine translation (Graves et al. 2013; Bahdanau et al. 2014; Cho et al. 2014).

The global view of the architecture is shown in Fig. 21. Initial natural language input is embedded using a bidirectional LSTM (Fig. 22) in the

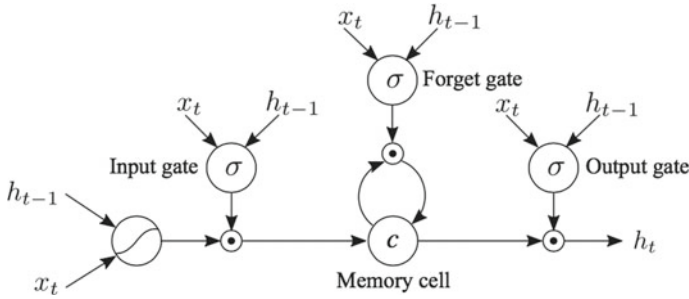


Fig. 22 Long Short-Term Memory (LSTM) unit

encoder unit. We then concatenate a one-hot embedding of the input words to the memory cell vectors of the LSTM using a linear layer. This is done by the aligner unit. Alignment allows the model to focus on parts of natural language instruction relevant to the current action. Alignment was proven effective in the context of machine translation and machine vision (Bahdanau et al. 2014; Volodymyr et al. 2014; Ba et al. 2014; Hu et al. 2015). The model learns to align based not only on the high-level input abstraction, but also on the low-level representation of the input instruction, which improves performance. The aligner is simply a linear layer $z_t = \sum_j \alpha_{tj} \begin{pmatrix} x_j \\ h_j \end{pmatrix}$, where h is the encoder embedding, and x is the original input. The weight α is computed by a nonlinear function.

$$\beta_{tj} = v^T \tanh(Ws_{t-1} + Ux_j + Vh_j)$$

$$\alpha_{tj} = \exp(\beta_{tj}) / \sum_j \exp(\beta_{tj})$$

where s is the hidden state of the decoder, and W, V, U are learned parameter matrices.

The final memory state of the encoder is used to initialize the decoder. This is another LSTM unit, which translates a sequence of world state observations into a sequence of robot actions. The robot takes those actions in real time, which affects the future world observations. Eventually a stop action terminates the output sequence. In order to choose which action to take, we pass the decoder output through another squashing nonlinearity and apply the soft-max layer to obtain a distribution over possible actions. Formally the decoder is another LSTM-RNN, which uses the world state representation y_t , the context of the natural language instruction z_t (produced by the encoder) and the decoder's previous memory s_{t-1} (aligning it with the

instruction context). We use a deep output layer (Pascanu et al. 2013) to produce a conditional probability distribution over possible actions to take next $P_{a,t} = P(a_t|a_{1:t-1}, y_t, x_{1:N})$.

$$\begin{pmatrix} i_t^d \\ f_t^d \\ o_t^d \\ g_t^d \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} T^d \begin{pmatrix} Ey_t \\ s_{t-1} \\ z_t \end{pmatrix}$$

$$c_t^d = f_t^d \odot c_{t-1}^d + i_t^d \odot g_t^d$$

$$s_t = o_t^d \odot \tanh(c_t^d)$$

$$q_t = L_o(Ey_t + L_s s_t + L_z z_t)$$

$$P_{a,t} = \text{softmax}(q_t)$$

In the equations above, E is an embedding matrix, L_o, L_s, L_z are learned parameters.

For training we use negative log-likelihood of the given action at each time step as the loss function.

$$L = -\log P(a_t^*|y_t, x_{1:N})$$

Inference is done over an ensemble of randomly initialized models, which is shown to have a denoising effect in deep neural networks (Sutskever et al. 2014; Zaremba et al. 2014; Vinyals et al. 2015). At each time step we choose the maximum a posteriori action under the trained model.

This model can be easily implemented in Python using the PyTorch tensor computation library.

Now it's time to discuss the data set we need to generate in order to train the designed model. We are interested in learning the model mapping from corpora of training data of the form $(x^{(i)}, w^{(i)}, a^{(i)})$ where $x^{(i)}$ is the natural language input (i.e. the instruction given to the robot), $w^{(i)}$ is the sequence of environment observations, and $a^{(i)}$ is the desired sequence of actions that robot should take in the context of the world observations in order to correctly follow the given instruction. After the model has been trained, it is able to predict a sequence of actions $a^{(i)}$ given a natural language instruction $x^{(i)}$ and environment observations $w^{(i)}$.

In a typical approach to this problem we would design two data sets. First one would come from a simulator, which allows to test ideas quickly before training on a real-world data, and the second one would involve training the refined model on more realistic data. For the purpose of this case study,

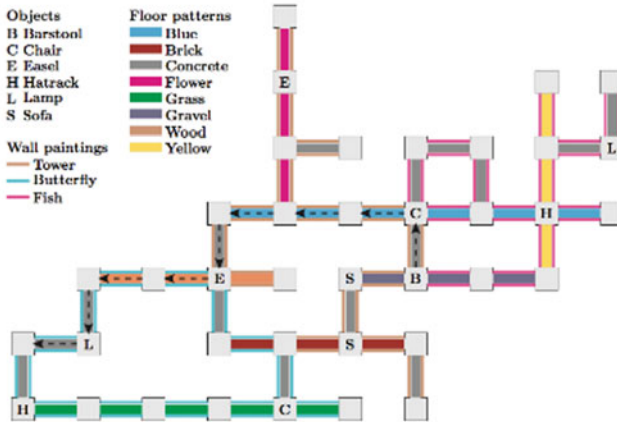
consider the first data set generated from an artificially generated environment (Fig. 23). The environment consists of a maze of hallways characterized by different texture patterns (grass, brick, wood, gravel, blue, flower, yellow, octagons). Various objects are placed within the halls (hat rack, lamp, chair, sofa, barstool, and easel) and the walls can contain several types of paintings (butterfly, fish, or Eiffel Tower). The natural language instructions come from humans who are told to navigate this world in a simulator. The instructions contain noise in form of ambiguities, spelling and grammatical errors, or being incorrect (e.g. using left when meaning right). Some sentences can not be mapped to any actions, and some produce unfeasible paths through the maze.

In our case the second data set would consist of Google Street View data from real-world environments (Fig. 24). We can write code to generate paths through different cities, paired with sequences of geolocation data, environment information (e.g. type of road), and sets of images covering 360 degree view at each location. Every path will be paired with natural language instructions, which include references to visible landmarks along the route. There is randomness in the choice of language, as well as landmark references. Because mentioned landmarks are detected using vision API (based on convolutional neural network models), it contains a degree of ambiguity, which makes it more realistic in comparison to the purely virtual dataset.

Having designed the model and generated the data, the next step involves training and testing the model on various objectives. Experiments are meant to test performance of the model, and suggest redesigns. The entire R&D process repeats until satisfactory results are obtained.

Once the model is trained to a satisfactory level. The next step would involve testing it on an actual car or a smaller robot operating in similar environment. Because we use Street View images, the images an actual car would see during a driving session would be similar to what it learned during training. Many other components are required to successfully deploy a physical system in a noisy environment such as real city streets. These involve engineering problems relating to actual steering and perception in real-time, which are beyond the scope of this case study.

The same type of model can be used to achieve many other objectives. For instance we might have a software agent instead of a vehicle, whose mission is to trade stock. In this case we could take this model, and replace the natural language input coming into the first encoder with news articles, Twitter feeds, as well as history of the stock market price variations of a selected group of indices over some period of time. In the decoder part, instead of feeding it images encoded by a convolutional neural network



Place your back against the wall of the “T” intersection. Go forward one segment to the intersection with the blue-tiled hall. This intersection [sic] contains a chair. Turn left. Go forward to the end of the hall. Turn left. Go forward one segment to the intersection with the wooden-floored hall. This intersection contains [sic] an easel. Turn right. Go forward two segments to the end of the hall. Turn left. Go forward one segment to the intersection containing the lamp. Turn right. Go forward one segment to the empty corner.

Fig. 23 Virtual world dataset



Fig. 24 Google Street View dataset

as vectors, we could feed in vectors representing real-time measurements of some particular stock index, or a group of related indices. The output from the model can then be composed of buy/hold/sell actions performed by the automated trading agent (instead of navigation actions). A new dataset would have to be collected composed of stock market histories in a sector of interest. Many other combinations of data sources are possible. For instance, we could use satellite images, together with geographically focused social media feeds, containing natural language conversations, as well as news articles, in order to estimate group characteristics of the people living in a given area, and their economic status, which can be used for credit estimation. This type of model augmentation and reuse in different settings is common in AI development. We would often take ideas that worked in one domain, and after adjustment apply them in another domain of interest.

References

- Abadi, Martin, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv preprint. [arXiv:1603.04467](https://arxiv.org/abs/1603.04467).
- Artzi, Yoav, and Luke Zettlemoyer. 2013. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics* 1: 49–62.
- Artzi, Yoav, Dipanjan Das, and Slav Petrov. 2014. Learning compact lexicons for CCG semantic parsing. In *EMNLP*, 1273–1283.
- Ba, Jimmy, Volodymyr Mnih, and Koray Kavukcuoglu. 2014. Multiple object recognition with visual attention. arXiv preprint. [arXiv:1412.7755](https://arxiv.org/abs/1412.7755).
- Bache, Kevin, and Moshe Lichman. 2013. UCI machine learning repository.
- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. arXiv preprint. [arXiv:1409.0473](https://arxiv.org/abs/1409.0473).
- Belkin, Mikhail, and Partha Niyogi. 2003. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* 15 (6): 1373–1396.
- Bengio, Yoshua. 2009. Learning deep architectures for AI. *Foundations and Trends® in Machine Learning* 2 (1): 1–127.
- Bengio, Yoshua, Patrice Simard, and Paolo Frasconi. 1994. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks* 5 (2): 157–166.
- Bengio, Yoshua, Benjamin Scellier, Olexa Bilaniuk, Joao Sacramento, and Walter Senn. 2016. Feedforward initialization for fast inference of deep generative networks is biologically plausible. arXiv preprint. [arXiv:1606.01651](https://arxiv.org/abs/1606.01651).
- Brand, Matthew. 2003. Charting a manifold. In *Advances in Neural Information Processing Systems*, 985–992.
- Bryson, Arthur E. 1961. A gradient method for optimizing multi-stage allocation processes. In *Proceedings of a Harvard Symposium on Digital Computers and Their Applications*.
- Bryson, Arthur E., and Walter F. Denham. 1962. A steepest-ascent method for solving optimum programming problems. *Journal of Applied Mechanics* 29 (2): 247–257.
- Cayton, Lawrence. 2005. Algorithms for manifold learning. In *University of California at San Diego Tech. Rep.*, 1–17.
- Chen, David L. 2012. Fast online lexicon learning for grounded language acquisition. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers*, vol. 14, 30–439. Association for Computational Linguistics.
- Chen, David L., and Raymond J. Mooney. 2011. Learning to interpret natural language navigation instructions from observations. In *AAAI*, vol. 2, 1–2.

- Chen, Xinlei, and C. Lawrence Zitnick. 2015. Mind's eye: A recurrent visual representation for image caption generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2422–2431.
- Cho, Kyunghyun, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. arXiv preprint. [arXiv:1406.1078](https://arxiv.org/abs/1406.1078).
- Chung, Junyoung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint. [arXiv:1412.3555](https://arxiv.org/abs/1412.3555).
- Chung, Junyoung, Caglar Gülçehre, Kyunghyun Cho, and Yoshua Bengio. 2015. Gated feedback recurrent neural networks. CoRR, abs/1502.02367.
- Collobert, Ronan, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12 (August): 2493–2537.
- Dai, Zihang, Zhilin Yang, Yiming Yang, William W. Cohen, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. arXiv preprint. [arXiv:1901.02860](https://arxiv.org/abs/1901.02860).
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint. [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- Donahue, Jeffrey, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2625–2634.
- Donoho, David L., and Carrie Grimes. 2003. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences* 100 (10): 5591–5596.
- Elman, Jeffrey L. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning* 7 (2–3): 195–225.
- Fefferman, Charles, Sanjoy Mitter, and Hariharan Narayanan. 2016. Testing the manifold hypothesis. *Journal of the American Mathematical Society* 29 (4): 983–1049.
- Fukushima, Kunihiko. 1979. Neural network model for a mechanism of pattern recognition unaffected by shift in position—Neocognitron. *Electronics and Communications in Japan* 62 (10): 11–18.
- Fukushima, Kunihiko. 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36 (4): 193–202.
- Gauss, Carl Friedrich. 1809. *Theoria motus corporum coelestium in sectionibus conicis solem ambientium auctore Carolo Friderico Gauss.* sumtibus Frid. Perthes et IH Besser.

- Gauss, Carl Friedrich. 1821. *Theoria combinationis observationum erroribus minimis obnoxiae*, werke 4.
- Gehring, Jonas, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. arXiv preprint. [arXiv:1705.03122](https://arxiv.org/abs/1705.03122).
- Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.
- Graves, Alex, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6645–6649. IEEE.
- Graves, Alex, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. arXiv preprint. [arXiv:1410.5401](https://arxiv.org/abs/1410.5401).
- Greff, Klaus, Rupesh Kumar Srivastava, Jan Koutník, Bas R. Steunebrink, and Jürgen Schmidhuber. 2015. LSTM: A search space odyssey. arXiv preprint. [arXiv:1503.04069](https://arxiv.org/abs/1503.04069).
- Hadamard, Jacques. 1908. *Mémoire sur le problème d'analyse relatif à l'équilibre des plaques élastiques encastrées*, vol. 33. Imprimerie nationale.
- Halevy, Alon, Peter Norvig, and Fernando Pereira. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems* 24 (2): 8–12.
- Harnad, Stevan. 1990. The symbol grounding problem. *Physica D: Nonlinear Phenomena* 42 (1–3): 335–346.
- Hebb, Donald Olding. 1949. *The organization of behavior: A neuropsychological approach*. New York: Wiley.
- Hemachandra, Sachithra, Felix Duvallat, Thomas M. Howard, Nicholas Roy, Anthony Stentz, and Matthew R. Walter. 2015. Learning models for following natural language directions in unknown environments. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 5608–5615. IEEE.
- Hinton, Geoffrey E., and Richard S. Zemel. 1994. Autoencoders, minimum description length, and helmholtz free energy. In *Advances in Neural Information Processing Systems*, 3–10.
- Hopfield, John J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences* 79 (8): 2554–2558.
- Hu, Kelvin, Jimmy L. Ba, Ryan Kiros, Kyunghyung Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. <http://proceedings.mlr.press/v37/xuc15.pdf>.
- Hubel, David H., and Torsten N. Wiesel. 1959. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology* 148 (3): 574–591.
- Hubel, David H., and Torsten N. Wiesel. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology* 160 (1): 106–154.

- Ivakhnenko, A.G. 1968. The group method of data handling—A rival of the method of stochastic approximation. *Soviet Automatic Control* 13 (3): 43–55.
- Ivakhnenko, A.G. 1971. Polynomial theory of complex systems. *IEEE Transactions on Systems, Man, and Cybernetics* 4: 364–378.
- Ivakhnenko, Aleksei Grigorevich, and Valentin Grigorevich Lapa. 1965. *Cybernetic predicting devices*. New York: CCM Information Corporation.
- Ivakhnenko, Aleksei Grigorevich, and Valentin Grigorevich Lapa. 1967. Cybernetics and forecasting techniques.
- Joseph, Roger David. 1961. *Contributions to perceptron theory*. Ithaca: Cornell University.
- Kalchbrenner, Nal, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. arXiv preprint. [arXiv:1404.2188](https://arxiv.org/abs/1404.2188).
- Karpathy, Andrej, and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3128–3137.
- Kim, Joohyun, and Raymond J. Mooney. 2012. Unsupervised PCFG induction for grounded language learning with highly ambiguous supervision. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 433–444. Association for Computational Linguistics.
- Kim, Joohyun, and Raymond J. Mooney. 2013. Adapting discriminative reranking to grounded language learning. In *ACL (1)*, 218–227.
- Kiros, Ryan, Ruslan Salakhutdinov, and Richard S. Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. arXiv preprint. [arXiv:1411.2539](https://arxiv.org/abs/1411.2539).
- Kollar, Thomas, Stefanie Tellex, Deb Roy, and Nicholas Roy. 2010. Toward understanding natural language directions. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 259–266. IEEE.
- Kong, Chen, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. 2014. What are you talking about? Text-to-image coreference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3558–3565.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 1097–1105.
- Kumar, Ankit, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, 1378–1387.
- LeCun, Yann. 1985. Une procédure d'apprentissage pour réseau a seuil asymétrique (a learning scheme for asymmetric threshold networks).
- Legendre, Adrien Marie. 1805. *Nouvelles méthodes pour la détermination des orbites des comètes*. Number 1. F. Didot.
- Lin, Min, Qiang Chen, and Shuicheng Yan. 2013. Network in network. arXiv preprint. [arXiv:1312.4400](https://arxiv.org/abs/1312.4400).

- Looks, Moshe, Marcello Herreshoff, DeLesley Hutchins, and Peter Norvig. 2017. Deep learning with dynamic computation graphs. arXiv preprint. [arXiv:1702.02181](https://arxiv.org/abs/1702.02181).
- MacMahon, Matt, Brian Stankiewicz, and Benjamin Kuipers. 2006. Walk the talk: Connecting language, knowledge, and action in route instructions. *Def 2* (6): 4.
- Mao, Junhua, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. 2014. Deep captioning with multimodal recurrent neural networks (m-RNN). arXiv preprint. [arXiv:1412.6632](https://arxiv.org/abs/1412.6632).
- Matuszek, Cynthia, Dieter Fox, and Karl Koscher. 2010. Following directions using statistical machine translation. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 251–258. IEEE.
- McCulloch, Warren S., and Walter Pitts. 1943. A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics* 5 (4): 115–133.
- Mei, Hongyuan, Mohit Bansal, and Matthew R. Walter. 2016. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *AAAI*, 2772–2778.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. arXiv preprint. [arXiv:1301.3781](https://arxiv.org/abs/1301.3781).
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, 3111–3119.
- Minsky, M. 1969. Papert. s., perceptrons.
- Mitchell, Tom M. 1997. Machine learning. 1997. *Burr Ridge, IL: McGraw Hill* 45 (37): 870–877.
- Mooney, Raymond J. 2008. Learning to connect language and perception. In *AAAI*, 1598–1601.
- Narayanan, Hariharan, and Partha Niyogi. 2009. On the sample complexity of learning smooth cuts on a manifold. In *COLT*.
- Neelakantan, Arvind, Quoc V. Le, and Ilya Sutskever. 2015. Neural programmer: Inducing latent programs with gradient descent. arXiv preprint. [arXiv:1511.04834](https://arxiv.org/abs/1511.04834).
- Olshausen, Bruno A., et al. 1996. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381 (6583): 607–609.
- Pascanu, Razvan, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2013. How to construct deep recurrent neural networks. arXiv preprint. [arXiv:1312.6026](https://arxiv.org/abs/1312.6026).
- Pennington, Jeffrey, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. *EMNLP* 14: 1532–1543.
- Peters, Matthew E., Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. arXiv preprint. [arXiv:1802.05365](https://arxiv.org/abs/1802.05365).

- Radford, Alec, Wu Jeffrey, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1 (8): 9.
- Reed, Scott, and Nando De Freitas. 2015. Neural programmer-interpreters. arXiv preprint. [arXiv:1511.06279](https://arxiv.org/abs/1511.06279).
- Roweis, Sam T., and Lawrence K. Saul. 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science* 290 (5500): 2323–2326.
- Ruder, Sebastian. 2016. An overview of gradient descent optimization algorithms. arXiv.
- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1985. Learning internal representations by error propagation. Technical report, DTIC Document.
- Rush, Alexander M., Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. arXiv preprint. [arXiv:1509.00685](https://arxiv.org/abs/1509.00685).
- Sacramento, Joao, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. 2017. Dendritic error backpropagation in deep cortical microcircuits. arXiv preprint. [arXiv:1801.00062](https://arxiv.org/abs/1801.00062).
- Sacramento, João, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. 2018. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems*, 8735–8746.
- Schmidhuber, Jürgen., and Sepp Hochreiter. 1997. Long short-term memory. *Neural Computation* 9: 1735–1780.
- Schölkopf, Bernhard, Alexander Smola, and Klaus-Robert Müller. 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation* 10 (5): 1299–1319.
- Sharif Razavian, Ali, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. 2014. CNN features off-the-shelf: An astounding baseline for recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 806–813.
- Shen, Yelong, Xiaodong He, Jianfeng Gao, Li Deng, and Grégoire Mesnil. 2014. Learning semantic representations using convolutional neural networks for web search. In *Proceedings of the 23rd International Conference on World Wide Web*, 373–374. ACM.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, 3104–3112.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement learning: An introduction*. Cambridge: MIT Press.
- Tellex, Stefanie A., Thomas Fleming Kollar, Steven R. Dickerson, Matthew R. Walter, Ashis Banerjee, Seth Teller, and Nicholas Roy. 2011. Understanding natural language commands for robotic navigation and mobile manipulation.
- Tenenbaum, Joshua B., Vin De Silva, and John C. Langford. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290 (5500): 2319–2323.

- Vapnik, Vladimir. 2013. *The nature of statistical learning theory*. New York: Springer Science & Business Media.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Viglione, S.S. 1970. 4 applications of pattern recognition technology. *Mathematics in Science and Engineering* 66: 115–162.
- Vinyals, Oriol, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3156–3164.
- Viola, Paul, and Michael Jones. 2001. Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2001. CVPR 2001*, vol. 1, 1–511. IEEE.
- Volodymyr, M., N. Heess, A. Graves, and K. Kavukcuoglu. 2014. Recurrent models of visual attention. NIPS.
- Wen, Tsung-Hsien, Milica Gasic, Nikola Mrksic, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. arXiv preprint. [arXiv:1508.01745](https://arxiv.org/abs/1508.01745).
- Werbos, Paul J. 1982. Applications of advances in nonlinear sensitivity analysis. In *System modeling and optimization*, 762–770. Cham: Springer.
- Weston, Jason, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. arXiv preprint. [arXiv:1410.3916](https://arxiv.org/abs/1410.3916).
- Weston, Jason, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merriënboer, Armand Joulin, and Tomas Mikolov. 2015. Towards AI-complete question answering: A set of prerequisite toy tasks. arXiv preprint. [arXiv:1502.05698](https://arxiv.org/abs/1502.05698).
- Williams, D.R.G.H.R., and G.E. Hinton. 1986. Learning representations by back-propagating errors. *Nature* 323: 533–536.
- Xiong, Caiming, Stephen Merity, and Richard Socher. Xiong. Dynamic memory networks for visual and textual question answering. arXiv, 1603.
- Yih, Wen-tau, Xiaodong He, and Christopher Meek. 2014. Semantic parsing for single-relation question answering. In *ACL (2)*, 643–648. Citeseer.
- Zaremba, Wojciech. 2015. An empirical exploration of recurrent network architectures.
- Zaremba, Wojciech, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. arXiv preprint. [arXiv:1409.2329](https://arxiv.org/abs/1409.2329).
- Zoph, Barret, and Quoc V. Le. 2016. Neural architecture search with reinforcement learning. arXiv preprint. [arXiv:1611.01578](https://arxiv.org/abs/1611.01578).



Use of Big Data in Developing Markets

Yi Huang and Wenlan Qian

1 Introduction

Big-data-based FinTech has witnessed spectacular growth in developing markets, as exemplified by the leading players in China such as the Ant Financial, Tencent, and JD Finance. While the FinTech industry has taken innovative approaches to promote financial inclusion by serving traditionally underbanked businesses and households, it remains an open question as to how they prevail especially in comparison with traditional banking. This chapter covers how businesses in developing markets use big data to solve moral hazard and adverse selection problems, increase the efficiency of matching borrowers and lenders, and address liquidity and other concerns.

FinTech comprises large technology firms whose primary business is e-commerce or social media. They have been entering banking in developing markets, which is considered a particular subset of broader FinTech innovations. In the beginning, FinTech firms' entry into banking started with

Y. Huang (✉)

Graduate Institute, Geneva, Switzerland

e-mail: yi.huang@graduateinstitute.ch

CEPR, London, UK

W. Qian

NUS Business School, National University of Singapore, Singapore, Singapore

e-mail: wenlan.qian@nus.edu.sg

payment services. Increasingly, they have expanded their businesses beyond payments into the provision of credit, insurance, and savings products, either directly or in partnership with incumbent financial institutions such as commercial banks (Frost et al. 2019; Buchak et al. 2017).

The main advantage of FinTech firms is their ability to exploit their existing trading networks and the massive quantities of data they generate. This business model quickly took off and has had astonishing growth in recent years. To date, big-data-based FinTech has profoundly changed the way consumers live and has reshaped the financial services industry in developing markets. Their huge success raises a number of important questions. For instance, do FinTech lenders leverage an informational advantage from their unique access to users' data? Or is their advantage technological due to methodological innovations such as machine learning, deep learning, and big-data algorithms that provide credit references and improve the modeling of credit risk? How has the development of FinTech stimulated economic growth? What are the other economic effects of and challenges faced by FinTech firms? What are the future trends in banking due to market competition with FinTech?

The rest of this chapter is organized as follows: In Sect. 2, we summarize how FinTech firms have used financial service instruments and big data to reduce the frictions in financial markets. Section 3 presents how FinTech stimulates economic growth in the aspects of facilitating small business growth, supporting entrepreneurial activities, promoting household spending that lead to financial inclusion and redistribution. We conclude in Sect. 4 with a brief summary and discuss the challenges faced in the era of FinTech and big data.

2 How Does Technology Reduce Financial Market Frictions?

2.1 Innovation in Financial Service Instruments

To overcome the underbanked issue, the developing world has recently turned to mobile banking after previously focusing on micro-finance. Aker and Mbiti (2010) and Mbiti and Weil (2011) show that a money transfer system via a mobile phone, such as M-Pesa in Kenya, improves individual outcomes by promoting banking and increasing transfers.

In addition, Allen et al. (2013) show that the Equity Bank in Kenya has pursued branching strategies distinct from other commercial banks; it has

targeted underserved areas and poor households. Equity Bank's presence has a positive and significant effect on households' adoption of bank accounts and bank credit, especially for Kenyans with low income and less education and for those who do not own houses. Also, the Equity Bank provides financial services to segments of the population that were typically ignored by traditional commercial banks. In the process of doing so, the bank has generated sustainable profits that show that it is a viable solution to the problem of financial access that has hindered the development of inclusive financial sectors in many African countries.

Furthermore, Gupta et al. (2019) study the positive productivity outcome from the expansion of mobile phone networks in rural India that has facilitated the adoption of high production of a variety of seeds and chemical fertilizers. In addition, they find that areas receiving mobile phone coverage experience higher adoption of these technologies.

The increasing adoption of financial service instruments also has helped to promote financial inclusion of the poor and underbanked areas. Higgins (2018) studies a natural experiment that created a shock to FinTech adoption; between 2009 and 2012, the Mexican government disbursed about one million debit cards to existing beneficiaries of its conditional cash transfer program. How did merchants respond to the consumers' adoption of debit cards? And how did it affect other consumers' FinTech adoption and consumption decisions? The results show that the adoption of debit cards had spillover effects on the adoption of financial technology on both sides of the market. Small retailers adopted point-of-sale (POS) terminals to accept card payments that subsequently led to more consumers adopting the cards. Two years after the government program, the number of POS terminals in treated localities had increased by 18% relative to untreated localities.

In addition to the studies that focus on practices and developments specific to one country, Demircuc-Kunt et al. (2018) conducted nationally representative surveys of more than 150,000 adults in over 140 economies. These surveys gathered information on the use of financial technology, including the use of mobile phones and the internet to conduct financial transactions. Their study shows that a total of 1.2 billion adults have obtained bank accounts since 2011, including 515 million since 2014. Between 2014 and 2017, the share of adults who had an account with a financial institution or through a mobile money service rose globally from 62 to 69%. In developing markets, the share rose from 54 to 63%. However, women in developing economies were 9 percentage points less likely than men to have a bank account.

Recent technological development has introduced innovation in screening and credit scoring by using machine learning, deep learning, and other

models based on predictive algorithms to improve the traditional credit risk model. Mishra et al. (2019) study the adoption of credit scoring in retail lending by private banks and state-owned public sector banks in India. The paper shows that there is a significant gap in the level of adoption of credit scoring technology between the introduction of the technology and two years later. For borrowers with prior relationships, public sector banks' inquiry rates are lower than those for their private peers who subject virtually all loans to a credit bureau check.

2.2 Using Big Data to Improve Credit Provision

In contrast to the government-promoted adoption of credit scoring, the private sector in China took the lead by using the big-data approach in providing banking services to solve moral hazard and adverse selection problems, increase the efficiency of matching borrowers and lenders, and to address liquidity and other concerns.

FinTech credit in China has seen a dramatic increase during the past few years. According to the annual surveys conducted by the *Cambridge Centre for Alternative finance* (Rau 2019), FinTech companies in China have become the largest market for big data lending in absolute terms. There were 2,525 online credit platforms by the end of June 2017. This was roughly 3% of the overall outstanding household credit in China. The average rate of return for investors rose to 9% with the maturity of lending decreasing to 9.5 months. Cumulatively, the industry had 76.23 million participants (investors and borrowers), or about 5.5% of the Chinese population.

Also, the in-house research institution of the Alibaba group, Luohan Academy (2019),¹ shows that digital technology greatly reduces adoption lags and lifts penetration rates for less developed countries and regions that thus enlarges growth opportunities for all. Despite many challenges, a new development paradigm is emerging; one that is inclusive, sustainable, and may not require massive investments up front. For example, as of July 2016, Ant Financial in China, the world's largest mobile and online platform, provided credit lines to more than 2.9 million firms or vendors on its e-commerce trading platform, Taobao.

What are the unique features of big-data-based FinTech lending compared to traditional banking? Does FinTech have an information advantage due to their access to users' data or from technological advantages that arise from

¹ Luohan Academy Annual Report (2019): Digital Technology and Inclusive Growth.

innovations such as machine learning algorithms to provide credit references and to improve the modeling of credit risk? Chen et al. (2019) show that one clear feature of FinTech is P2P financing, or crowdfunding, that is the collection and acquisition of personal information. E-commerce credit providers have access to a vast amount of information on their clients, such as e-commerce transaction data and online financial and behavioral data that include anonymized records of purchases, payments, online shopping, money transfers, wealth products, house ownership, rental information, and social networks. This information helps reduce the key challenges in traditional financial services—adverse selection and moral hazard problems due to information asymmetries (Stiglitz and Weiss 1981). The use of technology and big data make the provider's information collection much less costly and much more efficient especially for new or small businesses compared to traditional banks.

Another important feature of FinTech lending that distinguishes it from traditional banks is information processing and decision-making through big-data algorithms. FinTech and big data enable the better provision of credit by substituting subjective decisions (made by individuals) with automated decisions based on hard information (e.g., Liberti and Petersen 2019). By replacing soft information in relationship lending with hard information, FinTech makes loan processing much faster, cheaper, and more efficient.

Moreover, FinTech lending is more effective in both post-loan monitoring and debt enforcement. Traditional bank monitoring relies on public disclosure of firms' financial activities and covenant design, while FinTech lending can utilize real-time and high-frequency data based on the multidimensional metrics of borrowers. Lenders can more accurately evaluate the credit and other operational risks and thus prevent fraud and default (e.g., Fuster et al. 2018).

Huang (2019) studies how contract enforcement affects FinTech lending that is not collateral-based. Using data from Alibaba on online lending to SMEs, he finds that there is a negative relation between FinTech lending and contract enforcement in China. This significantly differs from traditional bank lending which is positively correlated with contract enforcement. Also, firms located in cities with poor contract enforcement are willing to use more FinTech lending and pay higher interest rates as they face more financial constraints in accessing bank credit.

Chen et al. (2019) also find that the enforcement procedures and strategies of FinTech firms depend on real-time models, and they are highly algorithmized. In fact, it is very difficult for borrowers of FinTech lending to default compared to borrowers of traditional bank loans. This is because FinTech

firms can “track” their addresses and locations (based on daily consumption records in the case of Taobao, for example), identify their related parties, and use other various ways to contact them. There are also implicit covenants for borrowers if they fail to repay the debt because of the potential sanctions by the FinTech lender. For example, FinTech lenders also provide platform services to businesses and have access to their payments for goods or their digital wallets, and as a result default is very costly for these businesses.

3 How Does FinTech Stimulate Economic Growth?

3.1 Mobile Payment and Small Business Growth

Thanks to the rapid development of smartphone and financial technology, mobile payment is another key growth area, especially among major emerging economies. For example, the total amount of mobile payments in China reached 42 trillion USD in 2018, 12 times as large as the value in 2014.²

Mobile payments can benefit businesses by reducing transaction costs for both consumers and businesses. Mobile payments cut operating costs for businesses that are particularly relevant for small business creation and growth. For example, mobile payments help small businesses to circumvent the cost of handling cash as well as its leakage due to theft, which erodes as much as 5% of the revenue of small businesses. It deters the entry of small businesses that could have made a positive profit absent such costs. Agarwal et al. (2020) investigate this channel using Singapore as a setting: in 2017, leading banks in Singapore introduced a mobile payment technology that facilitated consumers’ ability to transfer funds to both merchants and other consumers. They use the registry data on the universe of the firms created in Singapore and find that after the introduction of this mobile payment technology, business-to-consumer industries witnessed a higher rate of business creation of 8% per month relative to business-to-business industries. Consistent with the cost-saving argument, the growth in business creation is stronger for mom-and-pop stores and industries with higher costs in handling cash. To further establish causality, they show that tourism areas, which are dominated by foreign tourists and therefore have a lower penetration rate of (local) mobile payment users, see no increase in business creation after the event. Self-employed individuals receive a higher inflow of income from a larger

² Source: People’s Bank of China.

consumption growth during the post-event period. This income reflects the real effect of the technological advance on business owners.

Mobile payments can also improve efficiency by reducing the duration of transactions. This improved efficiency moves customer traffic more efficiently and raises customers' effective demand, mostly in shops involving small transactions. Agarwal et al. (2019b) approach this conjecture by using data on mobile payment transactions and bankcard spending from a representative sample of consumers from a leading bank in Singapore. They find that the 2017 mobile payment technology not only significantly promoted the use of mobile payments but also generated a significantly positive spillover effect on merchants' sales through card payments: small merchants experienced a monthly increase in card sales of 3.5% compared to large merchants after the technology shock, and the sales growth only existed for the merchants who generally received small-size payments. They also find that the sales growth was primarily driven by newly established stores and that merchants were more likely to be patronized by new consumers. These results indicate that the rise in payment efficiency stimulated consumers to explore new stores.

3.2 FinTech Credit and Entrepreneurial Development

As the world's largest online retailer and one of the world's largest internet companies,³ Alibaba enables third-party sellers in China to take their own businesses to the web. Alibaba estimated its China retail marketplaces Taobao and Tmall.com had "contributed to the creation of over 15 million job opportunities with more than 10 million active sellers as of 2015."⁴ This growth enables Alibaba to access the vast big data collected from 300 million registered shoppers and 20 million vendors.

By extracting information from big data, FinTech companies can mitigate information asymmetry and provide better credit access to borrowers who are particularly constrained or excluded from the credit market. These characteristics allow FinTech credit to reach micro, small, and medium-sized enterprises (MSMEs), which are more likely to be underserved by traditional banking. Does FinTech credit have a real impact on MSMEs? The consolidated dataset of e-commerce and FinTech credit from Ant Financial serves as a unique laboratory to investigate this question. Hau et al. (2019a) examine

³ As of October 2014, Alibaba surpassed Walmart as the world's largest retailer. See "Alibaba Is Now the Biggest Retailer in the World," *The Telegraph*, October 28, 2014.

⁴ See "Alibaba Affiliate Ant Financial Raises \$4.5 Billion in Largest Private Tech Funding Round," *Wall Street Journal*, April 25, 2016; "Alibaba Job Boom: Jack Ma Chats with Trump About How to Create 1 Million US Jobs over 5 Years," *CNBC*, January 9, 2017.

whether FinTech credit can promote entrepreneurial growth. To identify a causal effect, they exploit a discontinuity in the probability of credit approval at a particular threshold value of the internal credit score: for businesses just past the threshold, the probability of credit approval increases by 37%. They show that e-commerce vendors with first-time online credit access increase their sales growth by 9.2% in the month following credit eligibility, and the sales growth is entirely driven by credit-constrained firms that use the credit line. Importantly, they show that FinTech can mitigate the geographic disadvantages of an unfavorable firm location: the firms that face a lower supply of local credit or that are far from bank branches rely more on FinTech credit and have larger sales growth after the online approval of credit.

FinTech credit from Ant Financial also reduces sales volatility of e-commerce merchants, as found by Chen et al. (2019). They find that in the three months after access to FinTech credit, firms' monthly standard deviation in the weekly growth rate of total sales decrease by 5.5%. The effect exhibits a countercyclical pattern: the reduction in sales volatility is more pronounced during contractionary monetary policy regimes or in regions with lower economic growth. The reduction in sales volatility also has a real effect: in the three months after the access to FinTech credit, firms experience a reduction in the likelihood of bankruptcy of 15%.

These results reinforce the previously stated two advantages of FinTech credit over traditional banking especially for MSMEs. First, the use of technology and big data allow FinTech lending to better acquire and process information than traditional banking. Due to their shorter history for traditional lenders to effectively evaluate their credit risk, younger firms face a higher level of information asymmetry. Indeed, the effect of FinTech lending on sales and volatility is stronger among younger firms. Second, with the new technology in both monitoring and debt enforcement, FinTech lending circumvents a poor legal environment and weak contract enforcement. Due to the difficulty of debt collection, traditional lenders are less willing to lend to MSMEs in the areas with poor legal protection and weak contract enforcement. FinTech lending, on the other hand, could monitor borrowers using real-time and high-frequency data from the e-commerce platform, and directly adopt sanctions and enforcement. Accordingly, the authors find the effect of FinTech lending is stronger among firms in areas with poor legal protection and weak contract enforcement.

FinTech lending also influences the accumulation of customer capital among MSMEs. Using data from the e-platform Taobao, Huang et al. (2019) find that access to FinTech credit leads to an increase in service rating by 9% in the month following the access to credit. Repeatedly, they find evidence

that indicates FinTech credit can help mitigate the frictions associated with traditional banking: the effect of FinTech credit on customer capital accumulation is more profound for firms in business downturns, poor collateral conditions, and weak legal enforcement.

3.3 Digital Payment and Household Spending

Digital payment instruments can help reduce transaction costs and bring convenience to consumers. For example, QR (Quick Response) code payments are a digital payment instrument that has received much popularity in emerging countries like China and Singapore. This instrument allows two parties in a deal to fulfill the transaction by simply displaying and scanning the QR code on their mobile phones. Agarwal et al. (2020) use transaction-level account information on 250,000 consumers from a leading bank in Singapore to show that consumers in Singapore significantly decreased their amount of ATM cash withdrawals after the introduction of the mobile payment technology. The convenience of mobile payments encourages consumers to expand the geographic scope of their shopping activities and explore new stores (Agarwal et al. 2019).

On the other hand, mobile payment is a less salient method than cash. For example, cash payment involves the act of parting with hard-earned money that is a more painful and memorable experience than payments that do not involve a physical exchange of money. The subdued salience of mobile payment entails a risk for consumers to miss their budget targets and induce overspending. Agarwal et al. (2019a) investigate this conjecture by using the receipt-level data on customers of a large Indian supermarket chain. This information comprises the receipt amount, payment method, and details of items purchased, and anonymized consumer identification. They exploit the demonetization in India that exogenously forced cash-dependent consumers to switch to digital payments: in November 2016, the Indian government unexpectedly announced the invalidity of the Rs. 500 and Rs. 1,000 notes, which account for 86% of the existing currency in circulation. They find that after the demonetization, cash-dependent consumers decreased their use of cash but significantly increased their total spending. One explanation for this overspending is that consumers switched to more expensive items within a category and became less likely to use promotional offers. However, more evidence is needed on other types of spending and in other markets to comprehensively evaluate the overspending hypothesis.

3.4 Inclusive Growth and Redistributive Effects

Technology helps overcome information and operational frictions, a natural outcome of which is financial inclusion and a more balanced growth in the real economy (see the report by Luohan Academy, 2019, and the speech by Spence 2019⁵). Hau et al. (2019b) show that FinTech credit allows a more inclusive entrepreneurial growth by leveling the playing field for the previously underserved sector or regions. Similarly, the effect of mobile payment technology on business creation is more profound for mom-and-pop stores and in poorer communities (Agarwal et al. 2020). Since small firms and less developed communities are more sensitive to the operating cost of handling cash, their results show that mobile payment facilitates financial inclusion.

Another example of technology promoting inclusion is e-commerce. For traditional offline trades, the upfront investment of setting up storefronts in destination cities, along with the information asymmetry between buyers and sellers due to geographic distance impose a significant cost on inter-city trades. These frictions make goods produced in one city less accessible to consumers from other cities. This lack of access leads to inequality in living standards across regions, even within a country. E-commerce may help mitigate the inequality in spatial consumption for two reasons. First, it helps reduce the fixed cost of setting up a brick-and-mortar store that allows firms to reach consumers in cities that they otherwise would not serve. Second, online shopping platforms facilitate the acquisition of information by consumers on goods from distant locations. Fan et al. (2018) provide evidence consistent with these two features. Applying a gravity model for both online and offline trading flows between province pairs, they find that the distance elasticity for online trading is only about one-third of that for offline trading, which indicates online trading is less constrained by geographic distance than offline trading. They also use online sales data at the city-category level from Taobao and find that residents from cities with smaller population sizes and market potentials spend a larger share of their expenditure online and enjoy a larger welfare gain from e-commerce.

However, the redistributive effect of technology hinges on consumers' ability or willingness to adopt technology. Couture et al. (2020) exploit a randomized control trial that facilitates rural parcel delivery and pickup of a large e-commerce platform in China to show how e-commerce influences consumer welfare. The detailed household survey data that cover demographic and spending information allow the authors to explore the

⁵ Speech by Michael Spence (2019): The End of Shareholder Primacy?

heterogeneous effect of e-commerce. They find that overall, rural households experience a reduction in the cost of living due to online purchases, especially for durable spending and for remote rural residents. However, this welfare gain is driven by a small group of consumers (15% of the sample) who are on average richer and younger than non-adopters. Thus, whether e-commerce can benefit the population as a whole remains an open question.

There are also debates on whether technology induces redistribution in such a way that hurts certain population groups or sectors in the economy. One such debate centers on whether e-commerce can hurt traditional retailers in brick-and-mortar stores. The evidence so far is mixed. Some studies provide evidence that e-commerce crowds out offline stores. For example, Chava et al. (2018) investigate this question by exploiting the (staggered) introduction of a large US e-commerce retailer's fulfillment centers into a county. Using sales and employment data for 3.2 million stores, they find local offline retail stores experience a reduction in sales (employment) by about 3% (2.5%). Brick-and-mortar retail stores also show a decrease in entry and an increase in exits, with the exit concentrated in smaller and younger stores. Using a payroll dataset for 2.6 million retail sector workers, the authors also show that the income of workers in the local retail sector drops by 2.4% after the increased accessibility to e-commerce. The effect is mainly driven by a decrease in the number of hours worked.

While the US evidence indicates direct competition between e-commerce and offline retailers and a significant crowd-out effect for the latter, China's experience is somewhat more positive. In particular, e-commerce and offline retail stores enjoy complementarities. For instance, online retailers can hardly offer the shopping experience provided by offline stores. The information consumers acquire from the internet may also reduce the search cost of offline shopping. Chen and Qian (2019) use the world's largest e-commerce shopping holiday, China's Singles Day, to show how large-scale online promotional events influence the sales of brick-and-mortar stores. Using debit and credit card transactions of a large, representative sample of consumers from the country's leading service provider of card payments, they find that consumer card spending at offline stores increases on Singles Day especially among products that depend on the shopping experience. Importantly, the increase in offline sales not only exists among large, geographically close, and established stores but is equally large among small, geographically remote, and new ones for which consumers face higher search frictions. This finding indicates spillover benefits of e-commerce to offline merchants.

4 Conclusion

Over the past few years, the innovation in finance technology has reduced the frictions in the financial markets, such as the payment system, bank account, credit reference system, and finally platform lending. The FinTech companies such as Alibaba provide automated credit to millions of underbanked small firms and households in rural and remote areas. In this chapter, we focus on the innovations and real effects of FinTech and big data in developing markets, which are a rapidly rising yet still nascent topic in the literature. Moreover, the FinTech credit in less developed countries is arguably the single source of credit for these small firms and households based on big-data algorithms, hard information, and automated loan process.

Compared to traditional banking, big-data-based FinTech lending has apparent advantages in better information collection, acquisition and accumulation, efficient loan processing, and comprehensive credit risk modeling both for ex-ante screening and ex-post monitoring. By replacing soft information with hard information and substituting human decisions with automated decisions based on numerical data, technology, and big data make information collection and loan decisions cheaper and more effective. Furthermore, FinTech lenders are more efficient and effective in loan monitoring and debt enforcement using real-time data based on multidimensional metrics, and the enforcement strategies are based on highly algorithmized models of credit risk.

Overall, the FinTech in developing markets has streamlined the financial service process and has improved resource allocation that in turn, has boosted business growth and household consumption. The financial technology is also an important driver of inclusive growth.

We would like to close this chapter by pointing to the imminent need to manage several challenges unique to the era of FinTech and big data. The intricacy comes in maintaining a good balance between maximizing the benefits of big data and at the same time delicately dealing with concerns over information monopoly, data privacy, and data inequality as well as the potential bias of the big-data algorithms and the dark side of machine learning. Researchers and practitioners have started to explore those potential issues (Fuster et al. 2018), but the answers are far from clear. Developing markets will remain at the forefront of the development of big-data-based FinTech from which we continue to learn in the years to come.

Acknowledgements We acknowledge the generous support of Luohan Academy and Ant Financial in providing data and information on the institutional details. We are grateful to Chong-En Bai, Patrick Bolton, Markus Brunnermeier, Long Chen,

Alex He, Bengt Holmstrom, QJ (Jun) Qian, Raghavendra Rau, Zixia Sheng, Neng Wang, Shangjin Wei, Wei Xiong, Liyan Yang and Feng Zhu for their insightful views and constructive comments.

References

- Agarwal, Sumit, Wenlan Qian, Yuan Ren, Hsin-Tien Tsai, and Bernard Y. Yeung. 2020. The real impact of FinTech: Evidence from mobile payment technology. Working Paper.
- Agarwal, Sumit, Pulak Ghosh, Jing Li, and Tianyue Ruan. 2019a. Digital payments induce over-spending: Evidence from the 2016 demonetization in India. Working Paper.
- Agarwal, Sumit, Wenlan Qian, Bernard Y. Yeung, and Xin Zou. 2019b. Mobile wallet and entrepreneurial growth. *AEA Papers and Proceedings* 109: 48–53.
- Aker, Jenny, and Isaac Mbiti. 2010. Mobile phones and economic development in Africa. *Journal of Economic Perspectives* 24: 207–232.
- Allen, Franklin, Elena Carletti, Robert Cull, Jun Qian, Lemma W. Senbet, and Patricio Valenzuela. 2013. Improving access to banking: Evidence from Kenya. World Bank Policy Research Working Paper No. 6593.
- Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2017. FinTech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130: 453–483.
- Chava, Sudheer, Alex Oettl, Manpreet Singh, and Linghang Zeng. 2018. The dark side of technological progress? Impact of e-commerce on employees at brick-and-mortar retailers. Working Paper.
- Chen, Haiqiang, and Wenlan Qian. 2019. Does e-commerce crowd in consumer spending offline? Working Paper.
- Chen, Tao, Yi Huang, Chen Lin, and Zixia Sheng. 2019. Finance and firm volatility: Evidence from Alibaba FinTech credit in China.
- Couture, Victor, Benjamin Faber, Yizhen Gu, and Lizhi Liu. 2020. Connecting the countryside via e-commerce: Evidence from China. *American Economic Review: Insights*.
- Demirguc-Kunt, Asli, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess. 2018. *Global Findex database 2017: Measuring financial inclusion and the FinTech Revolution*. Washington, DC: World Bank.
- Fan, Jingting, Lixin Tang, Weiming Zhu, and Ben Zou. 2018. The Alibaba effect: Spatial consumption inequality and the welfare gains from e-commerce. *Journal of International Economics* 114: 203–220.
- Frost, Jon, Leonardo Gambacorta, Yi Huang, Hyun Song Shin, and Pablo Zbinden. 2019. BigTech and the changing structure of financial intermediation. BIS Working Paper.

- Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery. 2018. The role of technology in mortgage lending. Working Paper.
- Gupta, Apoorv, Jacopo Ponticelli, and Andrea Tesei. 2019. Technology adoption and access to credit via mobile phones. Working paper.
- Hau, Harald, Yi Huang, Hongzhe Shan, and Zixia Sheng. 2019a. FinTech credit and entrepreneurial growth. Working Paper.
- Hau, Harald, Yi. Huang, Hongzhe Shan, and Zixia Sheng. 2019b. How FinTech enters China's credit market. *AEA Papers and Proceedings* 109: 60–64.
- Higgins, Sean. 2018. Financial technology adoption. Working Paper.
- Huang, Yadong. 2019. Does contract enforcement still matter? Evidence of FinTech lending in China. Working Paper.
- Huang, Yi, Chen Lin, Zixia Sheng, and Lai Wei. 2019. Digital finance and customer capital development. Working Paper.
- Liberti, José M., and Mitchell A. Petersen. 2019. Information: Hard and soft. *The Review of Corporate Finance Studies* 8 (1): 1–41.
- Mbiti, Isaac, and David N. Weil. 2011. Mobile banking: The impact of M-Pesa in Kenya. NBER Working Paper 17129.
- Mishra, Prachi, Nagpurnanand Prabhala, and Raghuram G. Rajan. 2019. The relationship dilemma: Organizational culture and the adoption of credit scoring technology in Indian banking. University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2019-54.
- Rau, P. Raghavendra. 2019. Law, trust, and the development of crowdfunding. Working Paper.
- Stiglitz, Joseph E., and Andrew Weiss. 1981. Credit rationing in markets with imperfect information. *American Economic Review* 71 (3): 393–410.



Use of Big Data in Insurance

Melanie R. N. King, Paul D. Timms, and Tzameret H. Rubin

1 Context

1.1 How Insurance Works

Modern Insurance was first borne out of the great fire of London, where leading city figures sought to introduce a mechanism to ‘safeguard the interests of the leading men of the city whose lives had been ruined by the destruction of homes, businesses and livelihoods’ (Boobier 2016, p. 2).

M. R. N. King (✉) · P. D. Timms

Wolfson School of Mechanical, Electrical and Manufacturing Engineering,
Loughborough University, Loughborough, UK
e-mail: m.r.n.king@lboro.ac.uk

P. D. Timms

e-mail: research@paultimms.co.uk

T. H. Rubin

School of Business and Economics, Loughborough University, Loughborough, UK

At its most fundamental, insurance a story of risk identification, quantification and mitigation. In this context,¹ risks can be *external* (i.e. the ‘insured risks’) or *internal* (Moral Risks, such as Fraud) (Huber 2002). Insurers must manage both these risk types successfully to keep a profitable business.

A customer (be that consumer or commercial), takes out an insurance policy for a set fee and duration in return for the insurer taking on the liability of an agreed set of risks (e.g. an accident occurs or some damage is caused). Should these agreed events then occur within the defined period, the customer makes a claim against the insurer, who is now responsible for the financial and administrative consequences as set out within the policy, and for returning the customer back to the position they were in before the risk occurred (a concept known as indemnity).

For example, should an insured property flood, the insurer would cover the costs to repair the property, alternate accommodation and any other foreseen cost agreed under the policy. The insurer will likely also have administrative commitments to coordinate the repairs and remove the burden from the customer.

There are three important concepts on how insurance works, which are pertinent to understand the impact of big data on the industry:

- **The use of proxy data to define risk probability:** Each insured (external) risk, according to the Insurance Risk Management Institute’s definition, has an *uncertainty arising from the possible occurrence of given events*. While typically only a small handful of ‘pure’ external risks exist for a given insurance policy, these pure risks cannot be directly quantified. Instead insurers use a range of proxy data as an indication to the likelihood of an insured risk occurring for a given policyholder. (For example, the risk of vehicle damage caused by a traffic accident is quantified by proxies such as car type, driver age and previous accident history.)
- **The use of risk distribution to manage exposure:** To manage a particular insurer’s exposure to risk (i.e. the cost of claims due to risks that have occurred) and by consequence their financial standing, risks are either pooled (to manage low-impact, higher frequency claims), or spread across multiple underwriters through reinsurance (where the risk is high-impact, but likely low frequency).

¹ The types of risk presented here are specifically about risks directly related to the procurement and enactment of insurance policies. Many additional risk types exist elsewhere, for example in the adoption of new technologies, in the operation of businesses, in the financing of insurance, or in business process change.

- **Information asymmetry:** Information asymmetry refers to transactions where one party has more or better information than the other. In current insurance contracts, the customer is the owner (or guardian) of the information needed by the insurer to successfully quantify the risks insured. Insurers are reliant on the truthful provision of information by the customer, and on the insurers ability to elicit correct and useful information. Particularly in the commercial insurance space, the majority of risks are estimated based on past data and/or using statistical models based on small samples, due to the sheer logistical and administrative challenge of collecting raw data.

1.2 Towards Transformation

The evolution of technology used in the financial sector is reshaping the broader landscape of linked financial services, including the more conservative insurance industry (Oliver Wyman 2017). Insurance companies are improving their business operations through digitalisation, exploring new business models, upgrading their propositions and developing innovative products and channels. Wyman (2017) cites five key technologies that are driving this change: Cloud computing, the Internet of Things (IoT), big data, artificial intelligence and blockchain.

1.2.1 Cloud Computing

Cloud computing refers to the storage, management and processing of data via a network of remote servers, instead of on a personal computer or physical server hardware located on company on-site premises. Cloud services extend the in-house capability of insurers and provide a flexible and scalable IT architecture to grow in size, pace of development and technical sophistication. Cloud service providers augment insurers capability to store, manage and process vast amounts of disparate data, through offerings at the physical *infrastructure* layer (Infrastructure as a Service—IaaS), the *platform* layer (Platform as a Service—PaaS) and the *software* layer (Software as a Service—SaaS). The cloud computing architecture provides insurers with a pick-and-mix IT environment to suit their digitisation vision and adoption strategy: from data centres, servers and storage solutions, high-performance computing, development tools, database and analytic environments to applications and cognitive computing services. The ecosystem of technologies and solution providers range from the dominant big tech giants such as

Microsoft, IBM, Google and Amazon to specific platform and software service providers specific to Insurance, such as Cognizant Cloud Insurance Platform (Cognizant 2019). However, not all proprietary solutions are private. In 2018 Allianz, one of the world's leading insurers set up an open source version of its Allianz Business System (ABS) so that anyone could contribute to value-added services, and easily develop systems on top of Allianz's platform ecosystem (Allianz Group 2018). In 2019 Allianz announced a strategic partnership with Microsoft (Allianz Group 2019), aiming to digitally transform the insurance industry by moving core pieces of its insurance platform to Microsoft's Azure cloud, with the aim of accelerating innovation by creating a plug-and-play architecture for larger insurers and for start-ups, often referred to as InsurTech companies.

1.2.2 Big Data and Analytics

The volume of information that insurers manage is huge. 'In 2012 the UK insurance industry created almost 90 million policies, which conservatively equates to somewhere around 900 million pages of policy documentation' (Boobier 2016, p. 5). Also, the quantity and array of different types of information available to insurers—*structured* in databases or *unstructured* in documents—from a range of sources is growing at an unprecedented level. 90% of the available data has been created in the last two years alone.

The term big data, to describe this phenomenon, was first popularised in 2005 by Roger Mougalas from O'Reilly Media, to describe a large set of data that is almost impossible to manage and process using traditional business intelligence tools (*The New York Times* 2013). The key to understanding the role and relevance of big data in relation to insurance is shown in Table 1, which lists the five commonly understood characteristics of big data.

Data analytics, or the analysis of data, is recognised as a key capability for insurers to unlock the value in their data. Sophisticated data analytics can garner valuable insights to enhance operational efficiency, business growth and manage risks. Boobier (2016) outlines a hierarchy of four types of analytics: *Descriptive* (reporting why, when and what), *predictive* (what will happen), *prescriptive* (what should be done about it) and next generation *cognitive* (analytics with real time, natural language and learning capability).

Insurance already has a legacy of analytical thinking as a result in the rise of Business Intelligence (to measure and control strategic and operational activity) and as a necessary result of actuarial processes. However, actuaries have found the use of classical statistics has several key limitations when it comes to big datasets and are increasingly looking towards machine

Table 1 Five characteristics of big data in the context of insurance

Volume	The quantity of information available; from records held in databases and data lakes; policy documents and images as digital files in network folders; and external datasets
Velocity	The speed at which the data is collected, made available and stored. From stable storage in data warehouses to real-time data in motion via monitoring devices, such as telematics from cars
Variety	Data comes in a variety of formats from many sources, in a combination of structured, semi-structured and unstructured. Semi-structured data presents problems as it is seldom consistent and difficult to aggregate. Unstructured data requires a significant amount of additional processes in order to make it usable for any sort of further analysis
Veracity	The reliability of the data. Not all data is equally reliable as it comes from different sources. One measure of veracity is the 'signal to noise' ratio of useful and usable information in the dataset. For example, the number of off-topic posts on social media used to gauge customers' reactions to a new media campaign
Value	The value of the data can be measured in different ways, for example, how valuable is the data in terms of giving deeper insight to a certain issue for example, the creditworthiness of a customer

Source Boobier (2016)

learning (ML) as an approach to solving complex problems, that simplifies goodness-of-fit calculations, protects against over-fitting and permits easy hyperparameter tuning. However, actuaries are finding that ML has its drawbacks: it is often much easier to apply judgement to statistic models and explain their findings to non-technical stakeholders (Labram 2019).

1.2.3 IoT and Telematics

The Internet of Things (IoT) is a commonly used term to refer to any device, vehicle, building and other products embedded with electronics, software, sensors and actuators along with network connectivity, which enables these objects to collect and exchange data. The amount of the world's data is predicted to grow from 33 trillion gigabytes in 2018 to 175 trillion by 2025, according to the IDC (Reinsel et al. 2018), and much of this will be produced by connected devices with embedded sensors.

Telematics is a form of advanced IoT application in vehicles and has been the focus of early IoT adoption in the car insurance. Commonly collected data include list of trips, trip duration and distance, average speed and driving style. Telematics can be collected from engine management systems in the vehicle, external devices that plug into interfaces in the car, smartphone apps

or driver dashboard applications. Telematics is used to assess and calculate risk, enabling new forms of car insurance, such as Usage-Based or Pay-As-You-Drive Insurance policies. A recent study proves the value of telematics-based data in the risk selection process with only three months of data being sufficient to obtain the best risk estimation (Baecke and Bocca 2017).

1.2.4 Artificial Intelligence (AI)

Defining the term artificial intelligence is fraught with controversy and many definitions tend to be either overly vague or overly scientific. We do not endeavour to define the term here, however, when financial institutions say they are using AI, they are typically referring to analytical automation based on machine learning (ML). ML extracts patterns from unlabelled data (unsupervised learning), or efficiently categorises data according to pre-existing definitions embodied in a labelled data set (supervised learning). ML systems rely on vast quantities of data to find hidden relationships and use reinforcement to improve its own performance automatically. The process relies on experimenting with a number of different algorithms and different ML approaches, on large sets of training data, in order to refine a ML model that can be used to perform specific pattern matching, categorisation or predictive tasks. There are two key requirements however, for ML to be applied on a wider scale. Firstly that there is enough compute power available to handle the amount of data that needs to be processed and the second, that there is enough data (Patel and Lincoln 2019). ML is different from previous forms of automation and statistical analysis because it enables many actions to be taken without explicit human instruction. The risk associated with using ML techniques therefore is of adequate oversight and understanding of how these black boxes work and the application of automated decisions based on algorithms can lead to the propagation of biases that can be difficult to identify and root out. Insufficient transparency can also lead to damaged trust in financial services. The benefits of adoption of AI and the risks of getting left behind create powerful incentives for insurers to collect every more personal and private data and implement AI solutions on a rapidly accelerated timeline, which presents huge risks to the Industry in itself (Patel and Lincoln 2019) as well as ethics, governance and policy challenges (Renda 2019).

1.2.5 Blockchain

Blockchain is a distributed ledger technology used to store static records and dynamic transaction data, allowing for a collective bookkeeping in an immutable ledger of digital events. A distributed ledger can only be updated by consensus of the majority of the participants in the system and once entered, data can never be erased. Blockchain is a secure and transparent means of conducting business transactions that provides a single source of truth that is updated in near real time. Many insurers are experimenting with distributed ledger technologies for fraud detection, risk prevention and smart contracting. However, research has shown that there are legal risks associated with distributed ledger which regulators need to grapple with, these being; data privacy, insider trading and market abuse (Tarr 2018).

1.3 Challenges for Insurers

All insurers strive to excel in three key elements of their business, that is; operational efficiency (through cost reduction, claims management and productivity), profitable growth (through customer acquisition and retention, cross-selling and upselling) and risk management (through capital efficiency and operational risk management) (Boobier 2016). Throughout this chapter real-world use cases, describing the application of these new technologies and data-driven approaches in the insurance value chain, are provided.

Section 2 highlights some of the developments in the use of big data for traditional insurance functions, in particular, being better able to detect fraud and provide greater personalisation for customers. Section 3 presents examples of where data and analytics are changing the basis of competition by enabling, not only improvements to core operations but to launch entirely new business models. Section 4 outlines the rise of the InsurTech, because although insurers have valuable historic data, they may struggle to keep pace when competing with new start-ups who will often create business models on real-time new data extracted from IoT, social media and other digital records. It is important to note that the insurance value chain is embracing data-service providers who provide a particularly important capability in the ecosystem of products and services when it comes to big data. Finally, Sect. 5 presents a perspective on the ethical, trust and regulatory aspects of big data and artificial intelligence in the insurance industry.

2 Developments in the Use of Big Data for Traditional Insurance Functions

2.1 Leveraging New Forms of Data and Technologies

Insurers increasingly are employing advanced data analytics and artificial intelligence to help them make better decisions and improve business operations across the traditional insurance function chain, that is; pricing and underwriting; quote, bind, issue; policy administration and central systems; claims and settlement. Examples of application areas include smartphone sensor data to augment vehicle telematics, geospatial analytics and aerial imagery, drone inspections, virtual auto claims using integrated smartphone technology, robotic process automation to automate manual document workflow, workplace wearables and trackers, cyber risk analytics, hyperlocal weather analytics and conversational chatbots (CBInsights 2019).

Most significant area for improvement, however, is to be able to better detect fraud.

2.2 Being Better Able to Detect Fraud

According to the Association of British Insurers (ABI), in 2018, an average of 1,300 insurance scams were detected each day, with the average scam reaching £12,000 (Association of British Insurers (ABI) 2019). Given this is approximately £5.7 billion per year in detected fraud, the value of fraud detection to Insurance organisations is self-evident.

Insurance fraud can take many forms, with multiple opportunities across the Insurance Ecosystem. Fraud can occur internally or externally to an insurance organisation, at the point of underwriting or at the claims phase, and it can be Soft (opportunistic) or Hard (organised crime) (Viaene and Dedene 2004). The City of London Police IFED (Insurance Fraud Enforcement Department) presents a range of examples (City of London Police 2016) surmised below:

At the underwriting (policy) stage:

- Misrepresenting or not disclosing material facts to lower the insurance premium.
- Purchasing a policy using another's details to gain more favourable terms.

At the Claim's Stage:

- Claiming for damage that did not occur in the manner described.
- Deliberately causing an accident with an unsuspecting party.
- Including additional parties in a claim to increase pay-out.
- Two or more parties deliberately causing a collision and making a claim.
- Exaggeration of genuine injuries and/or the loss incurred as a result of a genuine incident, or of fictitious incidents.

Externally to the insurer:

- Illegal intermediaries (either an individual or a group) who misrepresent themselves as an insurance broker, agent or insurer for profit. In this scenario, either fake insurance policies are provided to the unsuspecting customers, or otherwise the illegal intermediary alters information provided by the customer to get a reduced policy cost from an insurer, siphoning off the difference and leaving the customer inadequately or not insured.
- Where associated professionals to the insurance industry (e.g. solicitors, engineers, doctors, vets, Accident Management Companies, etc.) provide false professional information to support a fictitious or enlarged claim.

Internally to the insurer:

- Employees within the insurance process submit fraudulent claims 'under the radar' due to their knowledge of the internal checks, or otherwise aid the process of fraudulent claims submitted by others.
- Sensitive information is obtained (data theft) from insurance company systems.

The ability to detect fraud brings significant benefits to insurance companies. In simple terms, possible fraud is identified through 'triggers' in the behaviour of and information provided by the customer (and through other available sources). Cases that have 'been triggered' then undergo further, more detailed investigation.

Due to the sensitive nature of insurance fraud, insurers are unsurprisingly coy about what triggers they use; we can however generalise that insurers are looking for outliers within the experience of their data trends. Fraud detection is inherently limited by the data sources available for the analysis.

It requires data matching and rules-based decision-making to flag cases for further investigation.

While insurers are, on the whole, considered effective in their ability to stop fraud, the exercise is a costly one. The nature of the analysis is high volume, prone to false positives and resource intensive. The general sentiment of the industry suggests that most types of insurance fraud are rising, in some cases exponentially, only adding to the pressure.

While new and improved sources of data present opportunity for fraud departments to better-identify fraudulent activity, existing methods alone will likely buckle under the quantity of information. To combat this, predictive analytics can be used to automate and augment current fraud defences.

Predictive analytics is the use of data, statistical algorithms and machine learning techniques to identify the likelihood of future outcomes based on historical data. While the general concept isn't new, recent advancements in pattern recognition (supported by the increase in data and computing capability) enable more accurate and efficient fraud detection.

Esure and Synectics Solutions: Precision Platform

Synectics Solutions Ltd., a software company in the data analysis sector, has developed a predictive analysis and data science platform entitled 'Precision' (Synectics Solutions Ltd., n.d.). It is a 'Software as a Service' (SaaS) offering aimed at connecting sources of risk data and enabling the advanced analysis of that data.

Esure, a major UK insurer for consumer lines such as motor, home and travel, is one such user of the Precision Platform. Being SaaS, it allows Esure to utilise advance analysis capability without the upfront investment in the underlying technology and the recruitment of upfront additional in-house data science expertise (however, as with the adoption of any new system or methodology, an amount of staff training is always required).

Based on the internal company data (fed from Esure into the Precision platform), integrated with data from SIRA (another Synectics Solutions service which aggregates a range of external risk data sources into a workflow-based offering), an array of algorithms and modelling techniques are used to evaluate the risks (consisting of a mixture of data science and machine learning methods), resulting in a risk score for each insurance claim. System-driven profiling alerts are triggered based on certain rules to support Esure's fraud detection workflows.

Esure combines manual referrals, cross-matching clear and suspect claims and policies and other intelligence with the Precision scoring and alerts to inform which suspected fraud activities are referred to the internal Special Investigation Unit. This predictive analytics service is used alongside traditional fraud trigger methods.

Adding this predictive component has increased the accuracy that Esure can achieve and, by recalibrating their decision model on the back of this new data

source, has enabled Esure to reduce the number of false positives referred to the Special Investigation Team.

2.3 Providing Greater Personalisation for Customers

New data on customers' habits, preferences, interests and abilities are generated every second online, tracked through cookies, advertising systems, social media platforms, through our purchase history and with the interactions we make with companies generally. Resultantly, a significant data profile now exists for each consumer looking to purchase any insurance line.

This influx of data, which can be bought and harvested by insurers, represents an opportunity to target, personalise and customise insurance policies, thanks to advances in consumer preference analytics. Opportunities are many-fold, depending on the level of change an insurance business may wish to adopt:

- **Targeted Marketing:** Perhaps the earliest use of consumer data in this list—online advertising has long enabled insurers to focus their advertising strategies on specific demographics, using a combination of social media data, search/viewing history and shopping habits. Insurers can utilise this data to support cross-selling and upselling with ever-increasing success rates.

More recently, user sentiment analysis can be used to understand which parts of an insurance policy (and reasons for buying the policy) will resonate most with the customer and adapt their marketing accordingly.

Knowing the right time to run an ad, send a message or offer a price can also have a big impact on its effect, which can be much better targeted given the massive increase in geospatial information available (for example, knowing when someone is travelling versus on the move to ensure that the customer is in a position to accept a policy offer). In a similar strategy, life events (e.g. the birth of a new-born child, moving home or change of job announcement) to trigger advertising of applicable insurance products, and flag changes which may affect existing policies.

- **Policy Price Optimisation:** Using a range of data but specifically including previous purchase history, insurers can now use prediction-based analytics to accurately and dynamically adjust the price for a specific insurance policy. This type of technology allows insurers to maximise their profit and reduce the amount of rejected policies, both for new policies and for renewals.
- **Customer relationship systems:** As part of a wider shift towards a customer-focused service provision, advanced customer relationship systems can now offer context-aware information which supports call handlers when fielding customer enquiries. For example, knowledge of ongoing claims, previous complaints or (for example) a recent wedding, can be flagged to the customer service representative in an easy to understand manner to help support a positive experience for the customer.

Furthermore, customer-centric systems can remove the friction of clunky, segregated business processes that can frustrate the customer (for example, the need to re-confirm the identity when changing departments or the time taken for the call handler to change systems to identify the previous complaint).

- **Policy recommendation systems:** Advances in customer service systems, and technologies including conversational AI, chatbots and recommender AI, has now reached a level of maturity whereby relatively accurate granular analysis of customer data, and response to questions, can guide the customer to specific products or services, and make recommendations on the appropriateness of a given policy.
- **Policy customisation systems:** The next logical step from recommendation systems, by accessing and understanding key data about each customer, insurers can now adapt what is included in a specific insurance policy to better suit a customer's requirements (or anticipated requirements), without a lengthy and resource-intensive onboarding process. As well as providing much more targeted cover based on the needs of the customer, the length of cover can be flexible, leading to on-demand insurance.
- **Telematics and Behaviour-driven pricing:** In a much more fundamental change in the working principles of insurance, telematic data and behaviour-driven pricing can be used to completely tailor insurance to the individual. Again, we discuss this more in Sect. 3 below.

Sunday Insurance

Sunday Insurance (www.easysunday.com), based in Thailand, is an InsurTech company that offers customisable consumer car insurance policies and business health insurance policies. It advertises itself as a “one size doesn’t fit all” company, Car insurance policies are based on provided customer data (note: not telematics).

The service is hosted on the cloud (in this case on Amazon’s AWS platform), which gives the company the flexibility to scale up with reduced outlay costs. Insurance premiums are calculated using artificial intelligence, specifically machine-learning -trained algorithms, which in this case are provided by H2O.ai.

The core proposition of Sunday is to make the insurance process simple for the customer and is therefore reliant of a fully functional website with few technical issues. While the insurance application still needs basic information about the customer (e.g. name, address, car type), Sunday Insurance is different in that, after selecting which baseline package you wish to start customising from (for example, Third Party vs comprehensive vs comprehensive with addons), it then allows individual customisation of various cover parameters.

While common add-on variables such as roadside assistance and windshield cover are available, the full range of additional options is much larger than a traditional insurer provides. This is in addition to the customisable cover limits.

It should be noted that, by transferring the choice of cover limit to the customer, it is reliant on the customer having a much-more informed understanding of the risks they need to mitigate and their resultant policy cover requirements. Less-informed customers may be encouraged to go for the cheapest option at the expense of having sufficient cover in the event of an accident.

Source Amazon Web Services, Inc. (2019).

3 New Opportunities from Big Data, Intelligent Applications and IoT

In Sect. 2, we presented new opportunities that broadly fall into the existing business practice of the insurance industry. However, there are also opportunities for completely new insurance concepts, driven in particular, by the rise of IoT connected devices. These opportunities can be exploited across the insurance value chain, by existing incumbents and newer InsurTechs alike.

An extension of the trend towards personalised insurance discussed above, new sensing devices can provide insurers with direct data for many insurance lines. This is a major change for the industry, as for the first time, insurers can access near-real-time data of assets in which to base policy premiums, either

retrospectively (based on data over a previous period) or reactively (based on the current state of the information received).

- In the motor vehicle sector, telematic data (transmitted from ‘black boxes’ or using customer mobile phone data) supply real-time location and speed data, which insurers can use as a much better proxy for driver behaviour than traditional markers (for example age or job description).
- A significant uptake in the use of wearable sensor devices (such as fitbit), or the fact that we many people keep their phones on their person all day, allows the tracking of fitness. Policies can be incentivised to encourage good health habits, further mitigating the overall risk of a claim over a long-term period.
- A rise in connected home devices can send a wealth of data to insurers, who could provide active security monitoring services.
- In the commercial sector, asset data (for example fridge temperature data) can be monitored, with insurance pay-outs automatically triggered based on certain conditions. Premiums can be reduced when customers can prove that their buildings have been operated within set criteria (e.g. that fridges have been maintained at 4 degrees and have not been switched off—which would cause food to spoil and a potential claim).
- In a different example, shipping cargo can be tracked across the world, and much more closely associated with the prevailing weather conditions, allowing for cheaper policies in good weather. If cargo is lost overboard, this can be identified much more quickly, and automatic pay-outs can be triggered.
- Predictive maintenance capabilities (for example on wind farms) enabled by connected sensors within the machines, can further reduce policy premiums, as serious failures can be identified and rectified before they occur.

3.1 IoT and Telematics for Customised Insurance Offerings

In what could be described as ‘level 1’ IoT data use, insurers can offer customised policies which are based on ‘real-world data’. The common denominator in these examples is that the insurers now have access to proxy data that is much closer to the real risks than before. This allows insurers to be much more granular in their risk management. Resultantly, policy premiums can be much more closely mapped to the actual risk, and the level of risk pooling can be reduced, if not eliminated.

While this will benefit some parties, this situation also presents a risk that, at the other end of this spectrum, insurance prices will increase further. As the averaging of risk (both in terms of granularity and in risk pooling) reduces, some customers may find themselves uninsurable.

Auto—Root Insurance

Root Insurance (www.jointroot.com) is a US usage-based automotive insurance company. Root Insurance tracks actual user behaviour using a smartphone application to evaluate the risk and provide a policy quote to the customer. At the expense of data privacy, Root Insurance says its customers receive rates up to 52% lower than for traditional insurance methods. A combination of data science and machine learning methods are used to identify key indicators of risk, such as movements like braking, route regularity, and the speed at which a customer turns a corner.

While some smartphone-based/telematics systems track behaviour over the lifetime of the policy, Root Insurance is different. Through a mobile phone application, a user's activities are tracked 24/7 for a two-week evaluation period (prior to the insurance policy starting). It is not possible to stop this tracking at any point during this evaluation period and still receive a quote.

After the evaluation period, a customised policy offer is provided, based on the data profile collected. After this point, no further data collected will affect the policy price. The customer can then switch off the tracking within the phone. If they fail to do so, data may still be collected, but it is only used to refine the insurance data models, rather than to influence policy price.

Using mobile phone sensor data (including accelerometer, gyroscope, global navigation system data and compass) presents a series of challenges to the insurer. Sensor data of this type is very noisy, of which risk evaluation models need to account for. Collecting this type of data also may not provide the necessary context in all situations—without additional information, the insurer cannot identify if you are the driver or the passenger, or if you are in a taxi. If the movement of a plane or train (at slower speeds) is akin to car speeds, this might also pick up without careful filtering of the data.

Sources Crunchbase, Inc. (2019) and CNBC LLC (2018).

Property—FloodFlash

FloodFlash (floodfash.co) is an InsurTech specifically focused on supplying building flood insurance cover using simplistic, parameter-based triggers. FloodFlash operates by providing customers with a proprietary sensor to install at their property. If water reaches a certain level, a cash 'settlement' is automatically triggered and paid out.

FloodFlash represents one way in which the global insurance protection gap can be lessened. High risk areas, where flood insurance would be prohibitively

expensive under traditional insurance models, can now be covered using parametric insurance, also known as ‘event-based insurance. Customers can reduce their premiums by adjusting the trigger level in which a pay-out is due and the amount that would be paid out in that event.

As only the water level verification is required and settlements are fixed, this represents a much less resource-intensive claims process. In this type of insurance arrangement, only cash is provided in the event of a claim (which is fully automated), meaning that the insurance company does not need to provide large swathes of administrative support typical of most insurers.

Source Oxbow Partners Ltd. (2018).

3.2 IoT, Behaviour-Based Insurance Pricing and Active Risk Mitigation

In more engaged (‘level 2’) IoT application examples, insurers can introduce insurance services which not only monitor user activity, but also seek to actively influence the behaviours of the customer(s) and the related risks. While there is some inherent influence on behaviour when customers provide real-time data (‘level 1’ applications discussed above), these ‘level 2’ applications are characterised by a more intrusive effect on the fundamental insured risks and the provision of an active service which reduces risk. To give two examples:

- In health insurance, insurers can encourage healthy lifestyles by incentivising exercise to ensure a longer-term reduction in health-related claims.
- In car insurance, insurers could potentially integrate with navigation systems and adjust insurance based on which route is chosen (e.g. a user could be presented with the choice of two routes and may choose the journey that takes longer but costs less to insure).
- In buildings insurance, a raft of new in-home CCTV cameras, smart speaker devices, thermostat controls and other connected technology can provide real-time data of the state of a particular dwelling or building. Insurers can monitor this data for security risks, building failures (for example, electricity or burst pipes) and engage with the customer to fix problems. For example, if the door is left ajar for a certain length of time, the customer could be alerted. Or if a smart speaker picks up the sound of breaking glass, they could alert the police to a potential intruder.

This active engagement with the customer to reduce the likelihood (or severity) of insured risks represents a fundamental change in insurance

mindset, by reducing the risk ‘upstream’ of the point of claim. By incentivising the customer to reduce the risk, both parties share benefits of successful risk mitigation. This arrangement is sometimes called Shared Value Insurance (Jais et al. 2017).

Health—Vitality

Vitality (www.vitality.co.uk) is a South African Insurance brand (owned by Discovery Holdings and used in many global markets) which provides health and life insurance policies. Vitality sells directly to customers, but also through various partnerships with other organisations/insurers in different markets.

The fundamental premise of Vitality is to reward customers for living a healthier lifestyle, thus reducing the risks to the insurer. Both health and life insurance products have a number of innovative features, but the major selling point for both policies is the use of the reward system.

The reward system tracks user activity data (e.g. walking, going to the gym, eating healthily or visiting the dentist) using information from devices like Fitbit or Apple watches. In exchange, customers can earn cashback, premium discounts and rewards (e.g. spa breaks, cheap flights, cinema tickets, free coffee or subsidised technology).

The available reward benefits are dependent on the ‘status’ of the customer. When first joining the vitality programme, customers are enrolled at the bronze level. As points are earned, the status can be increased to silver, gold or platinum, unlocking better rewards. At the premium level, rewards are significant (e.g. 40% off the premium price) to inspire customers to exercise more. The scheme also enables the purchase of highly desirable products (such as the apple watch at time of writing), based on reward points, at deeply reduced prices, attracting more customers and improving the data flow into the company.

This application uses both a carrot and stick approach, and should the level of exercise drop, the premiums go back up. This combined approach is suggested to change the fundamental human behaviours and psychology.

Smart Homes—Neos

Neos (neos.co.uk) is a smart technology and insurance provider (by partnering with Aviva). They provide (leak sensors, smoke alarm battery sensors, CCTV cameras) as part of the insurance policy offering. Customers pay more per month for more sensor devices but get additional peace of mind. The technology can also be purchased standalone.

The Neos offering is additionally a responsive one. Through an application, customers can manage their devices and also interact with Neos. Sensor devices are monitored (e.g. for security, for leaks), providing customer peace of mind. For example, if a water leak is detected, Neos can respond to that and make arrangements to investigate without the customer being at home.

4 InsurTechs and the Data Services Ecosystem

4.1 The Rise of InsurTech

InsurTech is a relatively new word that represents a subcategory concept from the broader financial technology world. It encompasses more than the narrow aggregation of activities combining traditional insurance with technology. It represents the ability to disrupt the incumbent value chain and allow new business models to form. InsurTech businesses often have the ability to position the customer at the centre of their offerings and provide customised services through improvement of information symmetry. The InsurTech definition therefore includes all innovation types: process, product, service and design, that can improve or create new insurance service and products.

Although InsurTech businesses are based on technology (and technological tools such as ML, AI, big data, IoT) it is not only their novel products and services, it is about the way in which these start-ups actually develop their products. Often these technology start-ups adopt agile principles and innovation practices, for example releasing Minimum Viable Products (MVPs) to the market in much quicker development and release cycles than larger firms are able to adopt. They are therefore more responsive to new developments in technology and external changes in the market. InsurTech innovations have been the subject of recent academic research. Stoeckli et al. (2018) have studied 208 InsurTech innovations and created an emergent model, classifying 14 transformational capabilities under five main themes: Digital infrastructure operations, data-driven infrastructure operations, digital service provisioning, insurance service development, customer network promotion and partner network promotion.

InsurTechs' ability to identify and link technologies to a specific service, quickly puts them ahead of the incumbents in many ways and is very much a customer-centric industry, focusing on how to improve and optimise customer experience. By improving the customer's experience and therefore relationship with respect to their insurer is a game changer in this industry. However, the first InsurTechs were not start-ups, but rather direct insurance companies that challenged existing incumbents in the retail motor sector, such as Geico in the US, Admiral and Compare The Market in the UK (Ricciardi 2018).

With the value of the InsurTech *modus operandi* being recognised by the major insurers, partnerships are forming to maximise the expertise and experience of both collaborators, big and small. In 2016 one of the biggest insurers in the UK, Admiral Insurance, partnered with a tech start-up Onezero Me

(onezero-me.com) to create a new data-driven car insurance product for young drivers. The ‘firstcarquote’ product was aimed at new drivers and used Facebook posts to analyse the personalities of car owners and set the price of their car insurance accordingly.

Facebook users who write in short and concise sentences, use lists and arrange to meet friends at a set time and place [...] would be identified as conscientious. In contrast, those who frequently use exclamation marks and phrases such as “always” or “never” rather than “maybe” could be overconfident. (Ruddick 2016)

However, it was not all plain sailing. The Guardian reported that Admiral pulled the product less than two hours before it was officially due to launch (Ruddick 2016) with concerns raised by both Facebook and the Open Rights Group. Jim Killock, executive director of Open Rights, said: ‘Young people may feel pushed into such schemes because of financial constraints. The right to keep things private shouldn’t be the preserve of those who can afford it’.

Christl (2017) highlights the extent of corporate surveillance in everyday life (noting that insurance companies were among the first to use statistical models to predict consumer behaviour) and the consequence of pervasive consumer surveillance. The report notes that in recent years a vast landscape of partially interconnected databases has emerged from various industries that ‘collect, analyse, acquire, share, trade, and utilize data on billions of people’. Furthermore, ‘much of corporate data collection and utilization happens invisibly, often with neither knowledge nor consent of the subjects’ (ibid., p. 5).

Many of these data providers are InsurTechs who provide business to business (B2B) data services specifically to collect and analyse vast datasets for insurers, however, not all provide data services about people.

4.2 InsurTechs That Provide Data Services to the Industry

4.2.1 Insurdata (US)

Insurdata (insurdata.io) is a US InsurTech firm that provides reinsurance underwriter provides services specifically with property-specific data to support their pricing, underwriting and portfolio management decisions.

Insurdata specialises in high-resolution, peril-specific exposures and building-level risk data, using technology that includes mobile augmented reality and 3-D model creation, providing both desktop and mobile solutions.

4.2.2 Betterview (US)

Betterview's (www.betterview.net) risk management platform helps carriers to identify and score roof condition and other property risks. Betterview does this by using machine learning and computer vision to analyse manned aircraft and satellite imagery. Betterview then combines this data with other sources such as building permits, assessors' data and historical weather to determine the characteristics and condition of a roof to determine the risk of future losses.

4.2.3 Rezatec (UK)

Rezatec (www.rezatec.com) provides a geospatial data analytics platform for the infrastructure, forestry and agribusiness sectors. Its proprietary algorithms process data inputs from satellite, airborne and ground sensors to predict outcomes for key assets across a range of industry verticals.

4.2.4 Cape Analytics (US)

Cape Analytics (www.capeanalytics.com) uses AI and geospatial imagery to provide instant property intelligence for buildings across the USA so insurers can more accurately assess a property's risk and value.

4.2.5 CrowdAI (US)

CrowdAI (www.crowdai.com) is an image recognition software that combines machine learning, computer vision and human intelligence to maximise value for self-driving car, automated drone and satellite image companies. Currently targeting enterprises, CrowdAI accurately and cheaply provides scalable, high-quality image annotation, with plans to enter industries like medical imaging and the military.

4.2.6 Groundspeed Analytics (US)

Within the property and casualty insurance market, massive amounts of data remain tied up in unstructured or semi-structured documents, such as PDFs and Excel files, to scanned images and emails. Unlocking this data presents a real opportunity for insurers, reinsurers and brokers to help improve underwriting efficiency and automating submission data capture, to identifying underwriting profit pools and market-wide changes in loss exposure. CBInsights (2019) reports that there is a notable investment by companies into commercial data automation, an activity at the early adoption stage in the Industry. Those, insurers who are investing in Natural Language Processing (NLP) for example are starting to unlock the value of unstructured and semi-structured data. For example, the commercial data automation firm Groundspeed Analytics (groundspeed.com) provides services specifically for insurance. They use artificial intelligence to harvest, normalise, enhance and extract data to client's specification so that it can be used to unlock the value in unstructured data.

5 Challenges

5.1 Ethics, Systemic Trust and Compliance

The explosion in available data that can be used as proxy indicators to insurance risk is shifting the balance of knowledge towards the insurance organisations and away from the customers. With so much information now available to insurers,² it is less of a question of what *can* be done and more one of what *should* be done.

Many consumers and small businesses do not trust the insurance industry. The Chartered Insurance Institute set out the challenge of trust in the industry (Chartered Insurance Institute 2018); with consumers finding the language confusing, having difficulty comparing policies and a poor understanding of products. Clearly insurers have a long way to go to win the trust of consumers who are inherently untrustworthy of both the insurance sector and new AI technologies.

While data itself brings a range of ethical risks, the greatest focus should be how the wider 'data system' (inclusive of people, processes and technology) makes use of this new data and the decisions made as a consequence of that

² Assuming that the insurers overcome the barriers to change and adopt these new information sources.

data use.³ It is this systemic perspective that defines whether it is ethical and trustworthy.

For example, a large majority of new data coming from the motor vehicle and the health insurance sectors is geolocate in nature. This is necessary to derive many of the key risk indicators (e.g. speed of travel, where a car is parked, how much movement a person does etc.). However, raw geolocation data allows insurers to see much more than this, if they were so inclined. In addition to driver behaviour triggers, the same source data could be used to identify the locations you have visited, who you have met, whether you have visited a hospital or if you are religious, among other data points.

As big data, AI and other related technologies are still new and emerging, global society is yet to reach a coherent view on what we should and shouldn't be doing with data. Clearly, unfettered and unscrupulous access to data is not a desired state for most, but insurers will need at least some access and trust to modernise their risk evaluation methods to utilise new data sources.

The successful adoption of ethical data use into society requires a multi-pronged approach. At one end of the spectrum, Governments will need to legislate for acceptable use of data. At the other, the insurance sector will need to push itself forward to adopt appropriate codes and behaviours that are compatible with the law and societies expectations.

5.1.1 It's Not One Ethics Perspective

Across the globe, the perspective on what is and isn't ethical data use changes. While perhaps an over-simplified view, we compare three countries on the subject of data protection:

- The European Union, with activities such as GDPR, is very much a people-first culture of data ethics that ensures the right for information self-determinism.
- In the USA, data protections come from consumer law, as opposed to any fundamental rights.
- China, which is known for active monitoring of its population, puts the rights for data access with the government over and above the individual.

Consequently, data systems designed and built in each of these countries will have fundamentally different principles of data protection built within

³ These decisions could be made by a person (the 'traditional' way of working), or in the future they could be made using AI technologies.

them. While the fourth industrial revolution is pushing society towards a connected world where borders are less relevant than before, insurers looking to adopt new data systems must be careful to ensure the systems they use are compatible with the local definition of data ethics.

It is likely that adopting a data system developed, for example in china, and then using it⁴ in the EU, without being explicitly designed to cater for the local ethical profile, will fall short of regulatory requirements and societal expectations. Even in good faith, emergent and unexpected differences in culture will likely introduce ethical incompatibilities.

The European Union's GDPR Regulations

In 2018, A General Data Protection Regulation (GDPR) was introduced which enshrines seven principles of personal data use into law (Information Commissioner's Office 2018).

1. Lawfulness, fairness and transparency
2. Purpose limitation
3. Data minimisation
4. Accuracy
5. Storage limitation
6. Integrity and confidentiality (security)
7. Accountability.

These principles guide and direct the trajectory in which businesses (in the EU and working with people in the EU) can collect and make use of personal data. The GDPR legislation is substantial, despite its fairly simplistic principles. Although the regulation has been active since 2018, many issues in its use are still yet to be resolved, and the insurance industry is still grappling with its correct application, as many anticipated industry guidelines are yet to be produced (Insurance Europe 2019).

Most significantly for businesses, the GDPR regulations have introduced rights to erasure (commonly known as 'the right to be forgotten') and rights to correction (when personally identifiable information held about a person is known to be wrong). Although noble in its intent, this undebatable right for the erasure and correction of personal data brings many issues for many insurance data systems, which were never built with this facility in mind. If this right is exerted, companies will likely find themselves manually searching, editing and removing data from their current systems.

This is not to say that newer systems, built on machine learning fair any easier. It is not yet clear how companies should handle erasure/correction requests where the previous data has been used within a machine learning model; would changing that data affect the performance of the model. Machine Learning methods are like baking a cake—you cannot get the eggs out once you've made it. If insurers end up having to re-run machine learning algorithms with corrected data, this will be of significant burden.

⁴ Or rather more likely, using data services that are based within another country.

Similarly, Blockchain, which by its intended design is an immutable and permanent record of transactions, runs headfirst into the problem of a right to erasure (Insurance Europe 2019). In many organisational implementations, it is impossible to correct data once entered. In others, while possible, needs the coordinated actions of multiple parties to 'force' a change across the distributed record, an incredibly resource hungry process.

The GDPR regulations are a good example of how data ethics vary by country. In the EU, individuals have the right for publicly available data to be removed. By contrast, the USA holds a much stronger freedom of expression bias, where it is thought that leaving information available for anyone to see is considered the fundamental right. Additionally, we can consider the US upcoming Clarifying Lawful Overseas Use of Data (CLOUD) regulations:

The USA's Cloud Regulations

The CLOUD Act is a federal law enacted in March 2018 that governs the handling of any data outside the United States. It is significant in its scope because it removes the boundaries so that it is irrelevant where the data is processed or stores.

In short, any US company, or company with a US presence, has to support the US authorities when it comes to any aspect of their jobs (including criminal investigations). It creates a legal framework in the US where the data a company holds, regardless of type or location globally, can be requested by US law enforcement. This could be personal or company data, from business details to trade secrets and intellectual property.

The US CLOUD Act is directly at odds with the EU GDPR regulations when it comes to personal data protections. Under GDPR, personal data can only be shared with other countries which apply the protections of GDPR and importantly, can only be done so where there is a mutual legal assistance agreement in place. Simply put, a request for personal data from a company covered by GDPR requires a court agreement to do so from the EU country involved. This is directly at odds with the CLOUD regulations, which require unfettered access without court involvement to that same data.

Consequentially, should insurers fall under both legal jurisdictions, they must inherently not comply with one set of the regulations, and risk significant legal recourse. It should be noted that this does not only apply to the insurers themselves, but also the companies that insurers use in the process of handling data. For example, Microsoft, Google and Amazon are all US companies widely used by UK insurers for cloud computing capability. Resultantly, any data processed by these companies (regardless of geo-location) is "fair game" under the CLOUD legislation.

Source Ionos (2019).

5.1.2 The Future Direction?

To understand further the future trajectory of data ethics, we look at the parallel work on ethics and trustworthiness in the context of AI. The use of AI and the use of data is inextricable in this systemic perspective; it is not an extreme position to suggest that AI is merely advanced data use, as modern AI is fundamentally a data-driven concept. While it is the 'bigger' societal impact of AI that is driving current ethics discussions, the consequence of this debate will ripple back through the supporting 'stack' of emerging data technologies, including big data.

A wide array of national and international bodies has recently published guiding principles on the ethical use of AI and data. While it is not the purpose of this book to review the detail of each and every guidance paper, a clear set of principles are being endorsed across multiple, independent organisations. To evidence three:

- The Organisation for Economic Co-operation and Development's (OECD) Council on Artificial Intelligence states that AI should be inclusive, sustainable, promote well-being, be human-centred, fair, transparent, explainable, robust, secure, safe and accountable (OECD 2019).
- The World Economic Forum states that data use should be secure, accountable, transparent, auditable, equal to all (fair and unbiased) and ethical (protecting the rights and aspirations of the vulnerable) (Hoffman et al. 2019).
- The IEEE Global Initiative for the design of Autonomous and Intelligent Systems states that systems should respect and protect human rights, increase human well-being, empower data agency, have evidenced effectiveness, are transparent, accountable, consciously guarded against misuse and used by competent persons.

While the above three examples use different words on occasion, the trajectory towards human-oriented, human first data use is clear. Of other notable trends, insurers must be aware of the trend towards transparent and accountable systems, which will require functionality not currently seen in many application cases. While these forward-looking principles are not yet enshrined in law, they are a clear indicator of intent. The OECD recommendations are pertinent, as these recommendations have subsequently been endorsed and adopted by the G20 group of countries (June 2019) (G20 2019).

While the IEEE is a standards body, we know that, over time, voluntary policies and adopted strategies become best practice, which in turn shifts the focus from ‘those who do’ apply good standards to ‘those who don’t’ as wider adoption spreads. Where commercial and societal pressures do not influence change, governments may then legislate to change voluntary standards into compulsory ones.

5.2 Regulation and RegTech

The reporting process for insurers is complex and can be very expensive. Regulatory disclosure provides the means to compare one organisation with another and is a legal obligation. Effective and accurate disclosure also provides evidence of an organisation’s ability to meet its obligations to its financial stakeholders, demonstrates solvency and that key stakeholders are adequately protected.

The UK has two regulatory bodies for the insurance industry; the Prudential Regulatory Authority (PRA) and the Financial Conduct Authority (FCA). The PRA promotes the safety and soundness of insurers and the protection of policyholders. The FCA regulates how these firms behave, as well as the integrity of the UK’s Financial markets. In recent years, following the growth of different business model within the insurance industry, data and risk management is no longer the mandate for the incumbent in this market but rather a set of data and data analytics suppliers that provide new mechanisms for risk management capabilities. Many of these suppliers are operating from outside of the UK domain. Thus, the regulators are constantly required to follow global trends and changes to make sure the general public good is preserved with respect to discrimination, pricing, data protection.

The regulator in Europe, to which UK Insurers need to comply, is the Insurance and Occupational Pensions Authority commonly known as ‘EIOPA’—a sort of industry watchdog. Solvency II (which came into effect in January 2016) is the European regulation that governs the amount of capital that an EU insurance company needs to remain solvent and whether their risk management practices are sufficient. The principle has been adopted by main other countries who have created their own regulatory models based on this such as the US’ Risk Management and Own Risk and Solvency Assessment (RMORSA). The new Solvency II regulation has come under much criticism. Firstly, the cost of compliance has been much higher than expected, with some major insurers incurring tens of millions of euros (Boobier 2016, p. 39; Butler and O’Brien 2019). Other criticisms include: the significant

computing power required to comply with reporting; the fact that sensitive data is leaving the premises, the country or both; and the overall skill needed for compliance in a smaller insurer makes it difficult to fund in-house. Increasingly insurers are looking to regulatory ‘service providers’ to support their compliance.

EIOPA is actively engaged with updating regulation in the Insurance sector while encouraging innovation at the same time. In a recent roundtable discussion of Industry experts, EIOPA acknowledged that in the context of rapid digitalisation, big data and AI: ‘regulatory and supervisory authorities have a role to play, by encouraging financial innovations while, at the same time, ensuring a well- functioning consumer protection framework and financial stability. In doing so it is necessary to respect key supervisory principles such as proportionality, market integrity and technological neutrality. Initiatives such as regulatory sandboxes, innovation hubs or public-private partnerships show that it is possible to be innovative in the approach to foster financial innovation’ (EIOPA 2017).

Over 50,000 regulations were published between 2009 and 2012 in the G20 (Butler and O’Brien 2019) and compliance is task that insurers are keen to digitalise where possible. RegTech is IT that helps firms manage their compliance and while some believe that RegTech will help determine the future of financial services by helping people innovate while generating trust from consumers (Roy 2019), others are much more sceptical.

While these technology systems offer powerful compliance tools, they also pose real perils. They permit computer programmers to interpret legal requirements; they mask the uncertainty of the very hazards with which policy makers are concerned; they skew decision-making through an “automation bias” that privileges personal self-interest over sound judgment; and their lack of transparency thwarts oversight and accountability. These phenomena played a critical role in the recent financial crisis. (Bamberger 2009)

Machine Learning, Natural Language Processing (NLP) technologies and the use of common ontologies have the potential to provide some real benefits for insurers and their requirement to comply with regulatory frameworks but, AI will do little to overcome the issue of human error in the so-called experts, systematic misconduct and fraud in the system (Bamberger 2009).

6 Conclusions

New technologies and the emergence of big data provide huge opportunities for new and existing organisations alike to make significant gains in the market, however, there are many external and internal influences that challenge organisations capability to maximise those opportunities. The challenge for incumbents is to simultaneously keep pace and exploit new technologies and data, react and adapt to disruptive InsurTech competitors in the market, while endeavouring to make sense of the often hyped and over-evangelised technological innovations on the market.

The application of digital technologies is a topic of vivid debate in the insurance industry and the proliferation of white papers, industry conferences and online knowledge bases alone can fuel this hype and act as ‘expectation generation devices’ for future insurance markets (Meyers and Hoyweghen 2018). The pressure to move quickly and deploy AI solutions to remain competitive, may lead to risks including insufficient testing and an over-reliance on AI and data specialists (Patel and Lincoln 2019). While AI adoption spurs competition, it could also lead to market concentration. Root Insurance Co (2017) uses AI to proactively target low-risk drivers, which presents ethical questions for the regulators as it has the potential to transform market structures. Shifting the burden of higher risk drivers to others.

While machine learning systems, leveraging the exponential growth of new forms of data, enables insurers to evaluate risks at a much more granular level, this capability will in fact disadvantage certain customers and challenge consumers’ concept of fairness. Customer trust could be damaged not only through insufficient transparency of algorithmic decisions, but the potential of AI models to propagate biases that can be difficult to identify and correct (Patel and Lincoln 2019). The concept of redlining (a historic discriminatory practice in property insurance, to deny credit or provide credit on poor terms for reasons unrelated to creditworthiness) has been around since the 1990s. However, recent research has raised concerns that ‘algorithms in the housing arena have the potential to operate in a manner that perpetuates previous eras of discrimination and segregation’ (Allen 2019, p. 219). An additional key risk for the insurance sector is to find and retain specialist talent who can design, develop, deploy, test and maintain AI systems—in order to mitigate the risk of ‘optimisation at the expense of social benefit’ (Patel and Lincoln 2019).

Findings from a 2019 (Capgemini and EFMA 2019) survey (of 75 insurance executives from across 20 markets) highlight that over 70% of insurers believe that: not only that advanced data management capabilities would be

critical to establishing a future-state insurance marketplace but also they need to shift focus to holistic risk solutions for customers.

Many insurers are now differentiating themselves by providing holistic risk solutions through value-added services. In the 'smart' home for example, these include security monitoring with links to emergency services, property concierge services and remote care monitoring and assistance focused on specific health and well-being needs. The most advanced of these service offerings are in France, where some insurers sell *télesurveillance* at a €200–€400 annual fee to more than 10% of their home policyholders: the most sophisticated are specialised value propositions for the elderly and a platform of property services delivered by a network of local certified providers (Carbone 2019). These new value propositions mark a step change in the industry: from selling insurance products to providing 'assurance services' and perhaps with it, a change in the trust relationship between customers and the industry. Customer perception of fairness and overall trust in the industry will remain critical in this wave of innovation.

References

- Allen, J.A. 2019. The color of algorithms: An analysis and proposed research agenda for deterring algorithmic redlining. *Fordham Urban Law Journal* 46 (2): 219–270.
- Allianz Group. 2018. Allianz sets up open platform [Online]. Available at: <https://www.allianz.com/en/press/news/business/insurance/180129-allianz-sets-up-open-platform.html> [Accessed 14 November 2019].
- Allianz Group. 2019. Allianz partners with Microsoft to digitally transform the insurance industry [Online]. Available at: https://www.allianz.com/en/press/news/business/insurance/191114_Allianz-partners-with-Microsoft.html [Accessed 14 November 2019].
- Amazon Web Services, Inc. 2019. Insurance case study [Online]. Available at: https://aws.amazon.com/solutions/case-studies/sunday_insurance/ [Accessed 14 November 2019].
- Association of British Insurers (ABI). 2019. Detected insurance frauds in 2018 [Online]. Available at: <https://www.abi.org.uk/news/news-articles/2019/08/detected-insurance-frauds-in-2018/> [Accessed 14 November 2019].
- Baecke, P., and L. Bocca. 2017. The value of vehicle telematics data in insurance risk selection processes. *Decision Support Systems* 98: 69–79.
- Bamberger, K.A. 2009. Technologies of compliance: Risk and regulation in a digital age. *Texas Law Review* 88: 669–739.
- Boobier, T. 2016. *Analytics for insurance: The real business of big data*. Chichester: Wiley.

- Butler, T., and L. O'Brien. 2019. Understanding RegTech for digital regulatory compliance. In *Disrupting finance. Palgrave studies in digital business & enabling technologies*, ed. T. Lynn, J. Mooney, P. Rosati, and M. Cummins, 85–102. Cham: Palgrave Pivot.
- Capgemini and EFMA. 2019. *World Insurtech report 2019*. s.l.: Capgemini & EFMA.
- Carbone, M. 2019. Smart home insurance strategy 101 [Online]. Available at: <https://www.linkedin.com/pulse/smart-home-insurance-strategy-101-mat-teo-carbone/> [Accessed 14 November 2019].
- CBInsights. 2019. *Emerging trends: What's next in P&C insurance*. s.l.: CBInsights.
- Chartered Insurance Institute. 2018. *Public trust in insurance: The challenge*. London: Chartered Insurance Institute.
- Christl, W. 2017. *Corporate surveillance in everyday life: How companies collect, combine, analyze, trade and use personal data on billions*. Vienna: Cracked Labs.
- City of London Police. 2016. Types of insurance fraud [Online]. Available at: <https://www.cityoflondon.police.uk/advice-and-support/fraud-and-economic-crime/ifed/Pages/Types-of-insurance-fraud.aspx> [Accessed 14 November 2019].
- CNBC LLC. 2018. A new kind of auto insurance technology can lead to lower premiums, but it tracks your every move [Online]. Available at: <https://www.cnbc.com/2018/10/05/new-kind-of-auto-insurance-can-be-cheaper-but-tracks-your-every-move.html> [Accessed 14 November 2019].
- Cognizant. 2019. Cognizant Cloud Insurance Platform [Online]. Available at: <https://www.cognizant.com/cloudinsuranceplatform> [Accessed 14 November 2019].
- Crunchbase, Inc. 2019. Ohio Unicorn Root Insurance Rolls On with \$350m in new cash [Online]. Available at: <https://news.crunchbase.com/news/ohio-unicorn-root-insurance-rolls-on-with-350m-in-new-cash/> [Accessed 14 November 2019].
- EIOPA. 2017. *EIOPA InsurTech roundtable: How technology and data are reshaping the insurance landscape*. s.l.: EIOPA.
- G20. 2019. G20 Ministerial Statement on Trade and Digital Economy [Online]. Available at: <https://www.mofa.go.jp/files/000486596.pdf> [Accessed 7 May 2021].
- Hoffman, W., A. Boral, and D. Olukoya. 2019. *Data collaboration for the common good: Enabling trust and innovation through public-private partnerships*. Geneva: World Economic Forum.
- Huber, M. 2002. Conceptualising insurance: Risk management under conditions of solvency. CARR Discussion Papers (DP 9). Centre for Analysis of Risk and Regulation at the London School of Economics and Political Science, London.
- Information Commissioner's Office. 2018. Guide to the General Data Protection Regulation (GDPR) [Online]. Available at: <https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/> [Accessed 14 November 2019].

- Insurance Europe. 2019. *Insurance Europe annual report 2018–2019*. Brussels: Insurance Europe.
- Ionos. 2019. *The controversial CLOUD Act: Effect on data protection and data security in the European Union*. Berlin: 1&1 IONOS Cloud GmbH.
- Jais, N., F. Lund, M. Pfitzer, and A. Rodriques. 2017. *Insuring shared value: How insurers gain competitive advantage by better addressing society's needs*. s.l.: FSG Inc.
- Labram, A., 2019. *Institute and Faculty of Actuaries: Statistics vs machine learning* [Online]. Available at: <https://www.actuaries.org.uk/news-and-insights/news/statistics-vs-machine-learning> [Accessed 14 November 2019].
- Meyers, G., and I.V. Hoyweghen. 2018. 'This could be our reality in the next five to ten years': A blogpost platform as an expectation generation device on the future of insurance markets. *Journal of Cultural Economy* 11 (2): 125–140.
- OECD. 2019. Recommendation of the Council on Artificial Intelligence. s.l.: OECD/LEGAL/0449.
- Oliver Wyman. 2017. *Technology-driven value generation in insurance: Industry report*. s.l.: Oliver Wyman.
- Oxbow Partners Ltd. 2018. *Insurtech-led change in the Lloyd's Market: A primer and strategic guiderails for managing agents*. London: Oxbow Partners.
- Patel, K., and M. Lincoln. 2019. *It's not magic: Weighing the risks of AI in financial services*. London: Centre for the Study of Financial Innovation (CSFI).
- Reinsel, D., J. Gantz, and J. Ryding. 2018. *The digitization of the world: From edge to core*. Framingham: IDC White Paper.
- Renda, A., 2019. *Artificial Intelligence: Ethics, governance and policy challenges*. Brussels: Centre for European Policy Studies (CEPS).
- Ricciardi, V. 2018. InsurTech definition as its own manifesto. In *The INSURTECH book: The insurance technology handbook for investors, entrepreneurs and Fintech visionaries*, ed. S. VanderLinden, M. Millie, N. Anderson, and S. Chishti, 6–8. s.l.: Wiley.
- Root Insurance Co. 2017. How Root prices car insurance [Online]. Available at: <https://www.joinroot.com/blog/how-root-prices-car-insurance/> [Accessed 14 November 2019].
- Roy, S., 2019. How RegTech could help determine the future of financial services. In *The RegTech book*, ed. J. Barberis, D.W. Arner, and R. P. Buckley. s.l.: Wiley.
- Ruddick, G. 2016. Facebook forces Admiral to pull plan to price car insurance based on posts [Online]. Available at: <https://www.theguardian.com/money/2016/nov/02/facebook-admiral-car-insurance-privacy-data> [Accessed 14 November 2019].
- Stoekli, E., C. Dremel, and F. Uebernickel. 2018. Exploring characteristics and transformational capabilities of InsurTech innovations to understand insurance value creation in a digital world. *Electronic Markets* 28 (3): 287–305.
- Synectics Solutions Ltd. n.d. Improving fraud detection rates through the use of predictive analytics [Online]. Available at: <https://www.synectics-solutions.com/Portals/0/pdf/Improving%20fraud%20through%20predictive%20analytics.pdf> [Accessed 14 November 2019].

- Tarr, J.-A. 2018. Distributed ledger technology, blockchain and insurance: Opportunities, risk, and challenges. *Insurance Law Journal* 29: 254–268.
- The New York Times*. 2013. The origins of ‘big data’: An etymological detective story [Online]. Available at: <https://bits.blogs.nytimes.com/2013/02/01/the-origins-of-big-data-an-etymological-detective-story/> [Accessed 14 November 2019].
- Viaene, S., and G. Dedene. 2004. Insurance fraud: Issues and challenges. *The Geneva Papers on Risk and Insurance* 29 (2): 313–333.



Technology and Working Capital Finance

Anil Savio Kavuri and Alistair Milne

1 Introduction

This chapter examines the use of technology in providing working capital finance to smaller enterprises. Technology is helping the provision of working capital finance through: (i) facilitating convenient online application with automated provision of supporting data; (ii) allowing better assessment, monitoring and management of credit and liquidity risks; and (iii) allowing a range of novel business models, including direct loan investment by individuals and investment funds, to compete alongside traditional bank provision of SME working capital.

What difference is this making? Our principal conclusion is that, while technology has the potential for substantially lowering the cost and increasing

A. S. Kavuri (✉) · A. Milne
Loughborough University, Loughborough, UK
e-mail: a.s.kavuri@lboro.ac.uk

A. Milne
e-mail: a.k.l.milne@lboro.ac.uk

A. S. Kavuri
Crawford School of Public Policy, Australian National University, Canberra, ACT,
Australia

the convenience and availability of working capital finance for smaller enterprises, we are still at the early stages of applying new technology to the provision of small business credit. It will be some years before the full impact of the new technology-based approaches is felt across the world and fully exploiting this opportunity will require far-sighted policy interventions.

While new technology business models for credit intermediation are important, there is little or no sign that these are coming anywhere close to supplanting the major role of banks in the provision of working capital to smaller businesses. The new technologies can and will be employed by banks as well as by non-bank intermediaries and platform lenders, so their impact is increasing competition in the supply of loans and supplying new or previously underserved market segments, not completely replacing existing arrangements. The principal exception appears to be for very small businesses, defined for example as those ‘micro-enterprises’ with less than 10 employees, where short-term high-margin technology-based credit by specialised lenders is helping address what has previously been very limited access to bank credit.

The predominance of banks as a source of working capital finance is a major source of systemic financial risk. One of the most dramatic developments during the global financial crisis in 2008–2009 was the collapse in bank provision of working capital finance, as the liquidity and solvency crisis affected banks across the world. A direct consequence of the contraction in the supply of trade finance was a fall in world trade of 12% in 2009.¹ Action since then by regulators under the Basel framework has forced banks to operate with substantially more capital and less reliance on short-term wholesale market funding. But addressing this systemic risk has also further limited bank supply of working capital finance. Technology-based credit does not yet appear to have matured to the point where non-bank provision can fully replace a systemic contraction in bank supply of working capital finance.

These judgements are reached from a relatively brief examination of a complex landscape. Working capital finance is provided under a great variety of national legal, regulatory and institutional arrangements. Working capital finance is itself a broad category that includes many subdivisions (invoice finance including factoring, purchase order finance, asset-based finance, international trade finance). Small firms obtaining working capital finance also differ vastly both in their financing requirements and in their risk-profile including the extent to which they can provide collateral as security against loans. Technology may lubricate but local and transaction-specific factors also play a major role in determining outcomes.

¹ Chor and Manova (2012).

Analysis is further hampered by fragmentary and partial data. Unlike for example, the syndicated loan or corporate bond markets accessed by larger companies, there are no comprehensive global data sources on SME and working capital finance. Researchers are forced to draw on partial data sets or case studies and must be cognizant of the many peculiarities of local markets to properly understand the supply of working capital finance.

This chapter can do no more than provide a broad picture of how technology is promoting the increased supply of working capital finance. It cannot go very far in quantifying this impact. It can though provide a 'road map', a guide that will allow practitioners, policymakers and researchers to see the overall landscape and therefore better understands how technology is creating new commercial opportunities, posing new challenges and questions for regulatory and business policy and throwing up new research questions.

The chapter is organised as follows. We first use the available data to paint a picture of the provision of working capital finance in various countries (Sect. 2). One striking fact that emerges here is that explicitly technology-based intermediation, for example the 'loan-based crowdfunding' that has come to prominence for example in China and in the UK, still represents only a small share of overall working capital finance. Banks and non-bank balance sheet lenders are also adopting new financial technologies in their more conventional established lending businesses. But this is difficult to quantify. There is no clean way to distinguish their use of financial technology in credit provision.

We go onto review the demand for different forms of working capital finance and the new emerging business models, supported by modern information and communication technologies, meeting this demand (Sect. 3). We then look at some prominent case studies of technologically enabled working capital finance, highlighting some successes but also dispelling some myths (Sect. 4). We conclude with a brief discussion of future prospects, suggesting that we are still at the early stages of a revolution that will bring us much closer to meeting the gap in financing the working capital requirements of smaller firms, but full exploitation of this gap has a long way still to go and requires a range of supporting policy measures.

2 The Provision of Small Business Lending in Various Countries

This section discusses the magnitude of lending to small and medium enterprises globally. Our principal finding is that, while fintech-based provision of working capital finance continues to grow rapidly around the world, its share of overall working capital finance varies substantially between countries and everywhere remains relatively small compared to traditional bank lending.

For reasons of data availability, we are unable here to distinguish working capital finance from other forms of debt finance to small and medium sized businesses. While clearly a limitation to our discussion, this does not overturn our conclusions. Aside from property development companies using their developments as collateral, small businesses have limited access to longer term debt. Most small business lending is working capital finance. Moreover, larger companies do not normally face problems in accessing working capital finance., having a range of options for external debt finance unless they are financially distressed. Hence the focus on this section on small business lending is appropriate for understanding the impact of technology on the supply of working capital finance.

Data provided by the Financial Stability Board provide a convenient starting point. A series of reports, of which the latest is FSB (2019b), collects a range of data on the magnitude and growth of bank and non-bank intermediation since the global financial crisis of 2008.² The Financial Stability Board also monitors SME lending.³ We use these sources to provide a bird's-eye perspective on the importance of technology-based lending in total credit intermediation.

FSB (2019b) reports a stock of global credit assets at the end 2017 of 169.5 trillion dollars (211% of global GDP), of which \$97.2 trillion are loan assets and \$78.9 trillion are bank loans. A separate FSB survey indicates that this bank lending approximately \$14.3 trillion or 18% of total global bank lending are to small and medium-sized enterprises.⁴ Noteworthy is the exceptionally high volume of small business lending in China. In 2017 the Chinese

² The motive for these FSB reports is monitoring the potential for destabilising 'runs' on non-bank institutions employing money-market funding to hold loans or other illiquid assets. This is a concern because these institutions typically are unable to access lender of last resort facilities from central banks. Their data collection is however broader, covering all non-bank investment in loans and other illiquid assets. Until October 2018 the FSB described such intermediation as 'shadow banking'. Recognising that such alternative forms of intermediation are not necessarily a source of systemic risk, they now simply refer to non-bank intermediation.

³ In FSB (2019a).

⁴ This \$14.3 trillion is our calculation based on Graph 7 in Appendix A of FSB (2019a) (graphing ratios to GDP from the FSB survey of SME lending for advance, emerging and developing countries)

economy amounted to 15% of global GDP, but bank lending to small and medium-sized businesses in China was \$7.4 trillion, some 61% of Chinese GDP and more than half of global small business lending.⁵ The same Financial Stability Board survey suggests that the US had a correspondingly small ratio of small business lending of only 1.5% of GDP or \$0.29 trillion.⁶

How important is alternative lending to small and medium-sized enterprises ('loan-based crowdfunding' to SMEs directly funded through lending platforms, i.e. peer-to-peer or marketplace lending) when compared to bank lending to SMEs? This is a difficult question to answer precisely. But an approximate answer can be given, showing that the stock of alternative finance lending to SMEs compared to total SME bank lending was still very small in 2017: around 2% in China and the UK; between 1 and 5.6% in the US depending on the assumptions behind the calculations⁷; and averaging less than 0.1% across all other countries.

These findings are obtained as follows. The various surveys from the Cambridge Centre for Alternative Finance of platform-based lenders (those that support loan investment without themselves having any loan participation) reveal a total global *flow* of new alternative SME lending in 2017 of \$103.4 billion.⁸ They provide no corresponding figures for the stock of alternative SME lending. However, illustrative calculation suggests that the ratio of stock to flow for end 2017 SME lending will be about 1.5 in countries such as the UK and China where growth has been most rapid (and about 2 in the US where growth has been slower).⁹ Using these estimated stock to

and figures for nominal GDP from the IMF datamapper. <https://www.imf.org/external/datamapper/NGDPD@WEO/OEMDC/ADVEC/WEOWORLD>.

⁵ This Chinese figure of \$7.4 trillion is computed in the same way as the figure in the previous footnote.

⁶ There are various institutional differences explaining this small US figure, inter alia: an important role for credit unions in lending to small business; widespread use of credit cards for funding business debt which may not be classified as bank SME lending; borrowing from finance companies for purchases of vehicles and other equipment; and securitisation of loans issued under the government SBA guarantee programme (Federal Reserve Board 2017).

⁷ We note that the Financial Stability Board reports a ratio of FinTech credit to bank SME credit of 25% for the US (FSB 2019a). They use the same sources as we use in our analysis from which we obtain an upper value of 5.6%. We are unable to make sense of their figure.

⁸ For the source of this figure see notes to our Table 1.

⁹ The assumptions behind this ratio of 1.5 are that the flow of loans has doubled every year for the past five years, that SME lending is evenly divided between maturities of 2, 3, 4 and 5 years with all loans are amortised on a straight-line basis. This yields the following tabulation of flows and stocks and the 1.5 stock-flow ratio. For the US, where growth in alternative lending has been slower in recent years, similar calculation yields a stock-flow ratio of about 2.

flow ratios, implies that the global stock of total alternative platform-based lending to SMEs at end 2017 was \$154.6, or 1.1% of global SME bank lending.

Comparison on the same basis reveals large differences between countries. In China and in the UK—the two countries where platform-based lending to SMEs has developed furthest—the end 2017 ratio of total alternative finance SME lending to total SME bank lending was 1.99 and 1.96%, respectively. The corresponding figure for the US was 0.99%. In the rest of the world excluding China, the US and the UK alternative finance lending was less than 0.1% of bank lending to SMEs.

The US has an especially large presence of non-bank balance sheet lenders providing credit to SMEs.¹⁰ Such non-bank balance sheet intermediaries (sometimes called ‘shadow banks’) lend from their own balance sheets, funded by debt and equity held by institutional investors. The Cambridge Centre for Alternative Finance reports US data for such non-bank balance lending to SMEs.¹¹ Including these along with platform-based lenders increases the US ratio of new technology-based SME platform lending to total Bank SME lending to 5.59%. If we allow though for the wider range of non-bank provision in the US and low volume of bank lending to SMEs, when compared to other countries, this suggests that the overall provision of working capital finance by technology-based lenders in the US is similar to the UK.

Despite its comparatively small size, alternative lending still represents the most prominent example of technology-based credit intermediation. As can be seen in Table 1, China is much the biggest market, followed by the US and then the UK where this alternative form of intermediation first developed. The other countries shown are the ten other countries with the highest

Total loan flow	100	200	400	800	160	3200
2 year loans flow	25	50	100	200	400	800
3 year loans flow	25	50	100	200	400	800
4 year loans flow	25	50	100	200	400	800
5 year loans flow	25	50	100	200	400	800
2 year loans stock	25	63	131	266	533	1066
3 year loans stock	25	67	144	296	598	1198
4 year loans stock	25	69	152	314	635	1276
5 year loans stock	25	70	156	325	660	1328
Total loan stock						4869

¹⁰ See Gawrych (2019), a recent webpage reviewing more than a dozen leading providers.

¹¹ In Ziegler et al. (2018).

Table 1 Alternative finance total and SME lending in the most important jurisdictions

Countries	Total alternative finance market in millions					SME lending in millions	SME as % of total alternative finance
	2013	2014	2015	2016	2017	2017	2017
China (\$)	5,560	24,240	102,000	243,000	358,000	97,430	27
US (\$)	4,400	11,560	28,400	34,530	42,810	1450	3
UK (£)	666	1,740	3,200	4,580	6,190	2039	33
Australia (\$)	27	104	398	610	1,149	23	2
Canada (\$)	44	86	207	335	868	9	1
France (€)	76	154	319	444	661	88	13
Germany (€)	65	140	249	322	595	71	12
Japan (\$)	87	115	351	398	349	189	54
Netherlands (€)	46	78	111	194	280	86	31
India (\$)	5	12	40	124	269	29	11
New Zealand (\$)	3	22	268	223	262	14	5
Italy (€)	1	8	32	127	241	24	10
Brazil (\$)	4	8	24	64	219	34	16

Note Alternative finance lending (also referred to as loan-based crowdfunding, P2P lending and sometimes as marketplace lending) where investment in loans takes place through an online platform without the platform participation. This includes unsecured personal lending, short-term property lending and SME lending

Source Authors calculations based on data from the database of the Cambridge Centre for Alternative Finance. We are grateful to Leyla Mammadov for her help with accessing this data

volumes of alternative finance lending. The share of SME lending in total alternative finance lending varies markedly. It is highest in the Netherlands (54% of the total) around one-third of the total in the UK, Japan and China and as low as 3% in the US.

This analysis reveals a growing but still relatively small contribution of technology-based credit to the supply of working capital finance. There are reasons for thinking that the impact is bigger than suggested by these aggregated figures:

- First, technology-based platform lenders are providing a relatively large share of lending the very smallest ‘micro-enterprises’. Cambridge Centre for Alternative Finance data indicate that in the UK for 2017 these accounted for some 30% of lending to the smallest companies with turnover of less than £2 million.¹²

¹² Reported in Zhang et al. (2018).

- Second, this analysis does not take any account of the adoption of technology-based tools of credit application and credit-risk assessment by established banks and finance companies. There is however no easy way to separate technology-based provision of working capital finance for these lenders.

As our case study analyses below suggest, the new providers of technology-based working capital finance have certain competitive advantages over established banks. They are not subject to the same high level of capital requirements as banks (though on the other hand they cannot benefit from any lowering of funding costs from bank deposit insurance schemes). Compared to banks they can often operate with relatively light requirements for compliance with 'know your customer' (KYC, effectively piggybacking on bank KYC compliance). They are not constrained by legacy systems and are hence able to offer a convenient and easy to use interface for business loan applications without for example any requirements for visiting branches. In some cases, they are able to use novel data sources for assessment of credit risks.

At the same time there may be limits to new business models. In both the UK and the US, alternative finance platforms have turned increasingly to institutional investment as a source of funding. In the US many non-bank providers operate through both alternative lending platform and balance sheet lending.¹³ Banks themselves will in time be able to develop better customer interfaces and employ novel methods of credit assessment. As we will discuss in the following section, banks will continue to have a comparative advantage in providing some forms of working capital finance.

3 Market Segments

This section discusses the various forms of working capital finance and the substantial differences in size and risk-profile of different small and medium-sized enterprises seeking access to working capital finance.

3.1 What Is Working Capital Finance?

Working capital as most narrowly defined is an accounting concept. In accounting terminology, gross working capital is total short term or current

¹³ See Ziegler et al. (2018).

assets, i.e. those that are expected to be turned into cash within one year. On the other side if the balance sheet are short term or current liabilities that are expected to be paid within one year. Net working capital, the difference between current assets and liabilities (expressed in absolute terms or as a ratio), is a measure of the liquidity position of the firm.

More broadly working capital can be used to refer to any investments with a relatively short and predictable repayment period, from expenditures such as wages, raw materials and inventories to purchase of equipment and marketing, that are expected to be paid off through generation of cash revenues in the short to medium term, i.e. around one to three years. In this broader sense working capital goes well beyond the narrow accounting concept, including also equipment such as computers and vehicles with shorter payoff periods and also off-balance sheet assets such as opportunities for future sales and predictable revenue opportunities of longer than one year.

Working capital finance is all forms of short-term and medium-term finance of working capital, paid back once the anticipated cash flows are realised. It thus covers all balance sheet liabilities except equity, long-term debt and hybrid debt-equity instruments.

Table 2 tabulates some of the main forms of working capital finance. This includes arrangements where assets are not owned by the enterprise. Consequently, assets and their financing are no longer recorded on balance sheet (invoice factoring and discounting; equipment leasing). Assets can be used as collateral for secured borrowing.

3.2 The Heterogeneity of Small and Medium-Sized Enterprises

A further challenge in studying the provision of working capital finance so smaller business is the heterogeneity of small and medium-size enterprises (SMEs). Definitions used for compiling SME statistics, usually based on number of employees or on annual revenue, vary substantially from one jurisdiction to another.¹⁴ At the same time a hugely diverse range of enterprises are classified as SMEs. Take the UK as an example. The standard definition is that a small enterprise is one with 50 or less employees, while a medium-sized enterprise is one with 51–250 employees.¹⁵ Enterprises with 10 or less employees are separately distinguished ('micro enterprises').¹⁶ Some banking

¹⁴ See FSB (2019a, section 2.1, p 7).

¹⁵ See Rhodes (2018).

¹⁶ Of the 5.7 million businesses in the UK in 2018 only one quarter had any employees at all, down from one-third in 2002. This shift to self-employment has driven in part by response tax

Table 2 The variety of working capital finance

	Definition and examples	Distinguishing features
Purchase order finance	Advance of money to the firm's supplier to meet a customer's order ^a	Unsecured with a repayment based on existing orders
Invoice finance/account receivables finance	Advance of funds to the firm by a third party in return for a percentage of the value of firm's invoices or a loan using invoices as collateral ^b There are different types	The factoring company has responsibility to follow up unpaid invoices from the firm's clients Unlike invoice factoring, the firm controls administration of debtors ^d
Unsecured credit	Invoice factoring/debt factoring. The firm sells to a third party the unpaid invoices from the firm ^d Invoice financing. The firm borrows money using invoices as collateral Bank overdrafts. A line of credit linked to the firms Line of credit. A separate financial that can be drawn on Term loans Credit cards	Secured loan. Repayment typically within around 90 days Unsecured. Based on prospective cash flows. Can be flexible, used as required by the firm or amortised on a fixed schedule
Other asset-based finance	Equipment finance. Loans of 3–5 years collateralised on the equipment Leasing Term loans secured on property	Secured. The vehicle, machine or property can be claimed if loan not repaid

	Definition and examples	Distinguishing features
Trade finance	Letters of credit forfaiting	Additional arrangements and instruments to deal with the challenges of international trade

^aSee Prosser (2015)

^bSee Business Expert (2019)

^cSee Prosser (2015) and Business Expert (2019)

^dSee Business Expert (2019) and Market Finance (2019)

statistics in the UK use revenue of less than £2 million as an alternative criteria for distinguishing micro-enterprises.

The available information on and the risks of credit provision are very different for microenterprises with less than 10 employees than for medium-sized enterprises with 50–249 employees. The smallest companies often have relatively undiversified revenues with the loss of a single customer or an essential employee can undermine their viability. Risk of fraudulent credit application is especially large for micro-business credit applications, when the business may not even exist in the first place or not have any reliable audited accounts.

Within these size categories financing needs vary hugely. An established business with strong and steady cash flows but limited opportunities for expansion is likely to have relatively little need for external working capital finance. External working capital financing requirements are much larger for newly established enterprises, when revenues are uncertain or have substantial seasonal variation, when margins are tight or when opportunities for expansion arise.

This great variety in both financing products, risks and financing needs, means that there are no simple standardised technology-based solutions for providing working capital finance to SMEs. Non-bank technology-based lenders have focused on specific niches where new approaches can give them a competitive edge over the standard established bank provision.

4 Case Studies and Examples

This section discusses some prominent case studies technology-driven change in working capital finance.¹⁷ We begin with the most prominent Chinese example, that of Ant Financial Services. We then look at a number of examples of balance sheet lenders in the US. We go on to describe two of the most prominent technology-enabled providers of working capital finance in the UK, MarketInvoice and Funding Circle. Finally, we review the role of technology in provision of international trade finance.

and administrative burden of formal employment, with employees of larger firms switching to self-employment to provide flexible service.

¹⁷ We have selected these case studies to illustrate the range of different approaches to technology-supported provision of working capital finance. There are many other examples that we could have used. In Argentina, Brazil, Mexico, Mercado Crédito provides working capital loans to entrepreneurs on Mercado Libre (Claessens et al. 2018). PayPal Working Capital has provided over \$2 billion loans as of 2017 since launching in 2013 in the US, UK and Australia (Dean 2017).

These case studies reinforce the findings of the previous sections. Each of these platforms operates in its own particular niche, providing specific working capital finance solutions among those listed in Table 2 to particular customer segments. While successful, none of these platforms are emerging as disruptors that are displacing the established role of banks in provision of working capital finance. To an important degree they are providing finance to smaller businesses that have been previously underserved by bank lenders.

4.1 Ant Financial Services¹⁸

Alibaba, the Chinese e-commerce conglomerate that is the leading provider of business to business e-commerce platforms (1688.com, Alibaba.com) auction platform (Taobao) and business to consumer sales online marketplace (Tmall). The company was restructured in 2014 as Ant Financial Services Group with minority ownership from Alibaba enabling it to raise external venture funding.

The growth of Ant Financial Services (AFS) has been substantial. The user base is a significant proportion of the Chinese population which provides growth opportunities. AFS is sometimes characterised as the most valuable FinTech in the world. There is no clear distinction between their business and consumer lending. AFS often provides loans to unincorporated businesses as personal loans.

AFS includes the hugely successful Alipay payments solution facilitating transactions domestically, both mobile to mobile payments and for transactions on Taobao and Tmall, and internationally. As well as wealth management services aimed at Alipay users, Ant Financial also provides a variety of technology-based forms of working capital finance. These include Zhao Cai Bao, launched in April 2014, is a marketplace lending facility embedded in Alipay that matches individual investors as lenders to individuals, micro-businesses and SME borrowers; and MYbank: a private online bank, established on 25 June 2015 with a mission to serve small and micro-enterprises.

MYbank relies on internet and machine learning in data collection and a big data computing in risk analysis. This allows MYbank to offer a highly convenient customer services, taking only 3 minutes to assess risks and decide the amount of lending and 1 second to issue the loan to an Alipay account, with no manual intervention during the whole process and flexibility of both timing of repayment and amount of borrowing, with the minimum be on 1

¹⁸ We are grateful to Chusu He for her help in writing this case study.

yuan. By April 2017 MYbank had loaned out 800 billion yuan (\$112 billion) far outstripping the loan volumes of any of the other technology-based credit providers around the world (though still a small share of the total of \$7.4 trillion of small business lending in China in 2017).

4.2 Balance Sheet Lending in the US

We now examine three prominent US balance lenders active in the technology-supported provision of working capital finance.

4.2.1 OnDeck

Founded in 2006, OnDeck is the most prominent of the technology-based lenders in the US, publicly listed as OnDeck Capital. In September 2018, it became the first non-bank online lender to exceed \$10 billion in cumulative loans originated to small businesses (OnDeck Capital 2018). It now has loaned out \$12 billion in total globally (OnDeck 2019a) to customers in 700 countries (OnDeck 2019b). In 2018, OnDeck had record number of originations of \$2.5 billion, 17% increase from 2017 (OnDeck 2019c). Originations have been over \$2 billion for each of the last three years.

On Deck provides terms loans between \$5,000 and \$500,000. Terms loans can range from 3 to 36 months, with weekly lines of credit. OnDeck's 2018 annual report provides a number of key financial and operating metrics. Over the past three years the loan yield—the rate of return that is achieved on loans outstanding and so effectively a measure of the average annual rate of interest on its lending—was 36.2% (2018), 33.8% (2017), 33.2% (2016).

During 2018, the average size of a term loan made by OnDeck was \$55,490 line of credit being \$33,689 (OnDeck 2019c). The typical customer is a 7-year-old business with \$450,000 gross annual revenue (compared with minimum requirements of a 1-year-old business with \$100,000 gross annual revenue). Thus OnDeck is very much operating in the microenterprise space, charging fairly high rates of interest for relatively short-term loans.

While focused on lines of credit and unsecured term loans, OnDeck has been preparing to move into asset-based finance. In 2019, OnDeck plans to begin making their own equipment loans ranging from \$5,000 to \$150,000 (OnDeck 2019c). OnDeck points out that the 'equipment finance market is ripe for disruption' aiming to target loans from \$5,000 to \$150,000 with 2–5 years secured by essential-use equipment.

4.2.2 Kabbage

Kabbage operates in a similar market to OnDeck providing convenient access to short-term credit for the smallest microbusinesses. The company provides credit up to \$250,000. Kabbage can qualify a loan of \$200,000 loan within only with access to business bank data. It can approve the loan within minutes by using its technology. Kabbage obtains ‘read-only access to banking information, marketplace information, shipment histories, accounting logs and web analytics’ (Bary 2019). A business needs to have revenues greater than \$50 K per annum and be at least one-year old. Businesses can choose between 6, 12 and 18 month terms (Kabbage 2019a).

Kabbage has now 200,000 customers (Kabbage 2019b) in 2015, the average line of credit was \$25,000, with average borrowing taking 7/8 loans a year totalling \$50,000 (Dahl 2015). By December 2019, Kabbage had loaned out in total more than \$6.5 billion since 2009, with \$2 billion in 2018 (Perez 2019). Kabbage finished the third quarter of 2019 with \$100 million in revenue, with loans originations of \$715 million, an increase of 43% year on year (Bary 2019).

The Kabbage business model differs from standard bank term loans, by allowing customers considerable flexibility in repayment. This is compensated by charging relatively high fees in the first months of a loan (so revenue is earned even if the loan is outstanding for only as short period). In 2015, CEO, Frohwein acknowledged that the average annual percentage rate of loans could be as high as 40%, with rates as high as 20% for the first 2 months, falling sharply in subsequent months (reported by Dahl 2015, their webpages provide similar illustrative calculations).

4.2.3 Biz2Credit

Biz2Credit is another leading US fintech-based business lender. It is a non-bank balance sheet lender. Founded in 2007, as of the time of writing Biz2Credit has arranged over \$2 billion in small business lending across the US. They also license their ‘AI-Powered Digital Banking’ platform Biz2X to other lenders including many commercial banks (a ‘software as a service’ SaaS platform that implements loan management, loan servicing, risk analytics and customer interface).

There is a short, informative, case study on the AWS website of the development of Biz2Credit <https://aws.amazon.com/solutions/case-studies/biz2credit/>. Biz2Credit originally worked as an ‘aggregator’, seeking to provide SME

loan applicants with an interface that would lead them to appropriate lenders, both banks and non-bank balance lenders.

It then found, however, that while this service worked well for loan origination, it did not provide the quality of 'end-to-end' services that they desired to provide for their customers. The total loan approval process could still take 7–10 days, because their client banks had to process the application through several different departments using different systems.

To offer the service they desired, Biz2Credit then made several developments in their business model:

- They began to also operate as a balance sheet lender, extending their platform to provide a unified operation that covered all the elements of the lending decision, including post loan servicing.
- Introducing technology for automating the transcription of information, e.g. from digital images.
- Developing a proprietary machine learning tool for credit risks assessment, using cash flows from banking records as well as standard accounting information, can more accurately assess credit risk. This though has yet to be tested in a major economic downturn.

This combination of a unified platform, together with automated information processing, allows them to make a credit decision and fund a loan in 24–48 hours. They are also though purely virtual, they provide their customers with telephone access to an adviser to help them get the right financing for their needs. Altogether they have provided arranged small loans to around 225,000 customers with a total value of \$2 billion, so their average loan size in the US is about \$10,000 (so on average serving even smaller businesses than OnDeck or Kabbage).

In May 2019 they announced the launch of a new service Biz2X, a 'Software as a Service' solution which allows banks worldwide to use the same unified platform that Biz2Credit has developed for their own lending. Their website <https://www.biz2x.com/index.php> indicates that their cloud-based platform can be integrated into existing bank systems and launched in 2 months (for smaller community banks) and 4 months (for global banks).

Biz2Credit has recently, in June 2019, successfully obtained \$52 million of Series B funding from WestBridge capital that will support their continued expansion.

4.2.4 Amazon Lending

Amazon Lending launched in 2012, provides loans to merchants on Amazon. Amazon Lending provides loans range from \$1,000 to \$750,000 for up to a year with annual interest of 6–17% (Dean 2017). In 2017, 20,000 SMEs in UK, US and Japan have taken loans (Dean 2017). Amazon investigates whether the SMEs are creditworthy and then invites them to apply for loans. Their technology can ascertain the financial viability of companies and whether they are having difficulties. Consequently, it can freeze accounts immediately (Dean 2017).

Note the unusual feature that these loans are invitation only: Amazon observes revenues and subject to minimum total sales of at least \$10,000 over the past year, good customer metrics and increasing or at least stable revenues they may offer a loan. Loan terms are for up to one year and at rates between 6 and 14%, rather lower than US credit card rates of 14–24%. The recipient can accept the loan, accept a smaller loan or reject. Once a loan is accepted then repayments are made automatically from the customer revenues passing through the Amazon site (in the event of non-payment it can also take further actions, e.g. confiscating inventory, preventing payments from the seller account).

The obvious disadvantage for borrowers is that, while the loans are pre-approved and so very fast to disburse, the Amazon loan offers are unlikely to be made at the time when finance is most needed.

Amazon Lending has subsequently expanded into the UK, Canada, China, France, Germany, India, Italy and Spain. In 2017 it reportedly loaned about \$1 billion with \$3 billion total lending since 2012 (though later press reports based on Amazon's annual reports lowered these initially announced figures). In 2017 it also announced plans for further expansion, something interpreted as a sign that the big tech players were making serious inroads into small business lending challenging established banks.

Recent reports by the Financial Times, in June of 2019, suggest that this expansion effort has not worked out as anticipated, with actual loan uptake disappointingly small. The volume of loans increased by only 4.7% in 2017 compared to 2016 and fell further to 2.6% in 2018.

Amazon publishes little information and there is no public data on loan losses from the programme. Some critical comments suggest that Amazon may not have effectively priced credit risk (though unclear that the problem was credit risk rather than simply lack of demand). The notion of a 'you wait until we offer you a loan product' does seem to have its limitations. But without adequate credit scoring, something that its customer data seems

as yet unable to support, it is difficult for it to operate more traditionally responding to loan requests.

Amazon is still suggesting that it may make a renewed effort to expand its lending business. Other large 'BigTech' players are more successful at credit. Most notably Ant Financial in China (see above). Paypal has a lending business that now loans more than \$1 billion per quarter. A more conventional approach based on traditional credit modelling may yet work for Amazon.

4.3 UK Technology-Supported Lending

- MarketInvoice.

New technology has been particularly prominent in providing invoice factoring and lending. Invoice factoring is when invoices are sold by businesses at a discount. In the event of late payment or non-payment, it is the factoring company, the financial company buying the invoices, that is responsible for collecting the amounts due. Invoice lending, when invoices are used as security against a short-term loan. In the event of late payment the business is responsible for collecting the due payments.

A UK alternative lender that has been particularly successful in this area is MarketInvoice <https://www.marketinvoice.com/>, providing credit against client invoices. Founded in 2011, it has already after eight years by the end of 2019 provided £2.7 billion of invoice and other business loans with 46,237 'trades' on their platform (<https://marketfinance.com/investor-statistics>). The gross interest rate charged has fallen gradually from 12% in 2014 to 7.5% in 2019 with loss ratios also falling to 1.6% in 2018. MarketInvoice investors must be high net worth, institutional investors or government authorities.

MarketInvoice supports two main products:

- confidential invoice financing secured against a business's entire portfolio of outstanding invoices. This involves an account migration, with the borrowing company setting up a new business account into which invoices are paid and from which market invoice is repaid, with an accompanying flexible line of credit related the total invoices due. The borrowing company must use an approved accounting software, effectively giving MarketInvoice oversight of its invoice book as well as direct claim on the invoice cash flows, but without passing on confidential client information to MarketInvoice are having to inform clients that their invoices
- selective (or single) invoice financing where an individual account is set up for the client to pay their invoices. Otherwise this is similar to confidential

invoice finance, although the request to transfer payment to a new account with a new bank could reveal that the invoice is being used as a cash flow commitment to repay borrowing.

They are also offering a purchase order finance solution and credit insurance and control services. MarketInvoice has been (but is no longer) a member of the UK Peer-to-Peer Finance Association, with their loan book financed by investment from institutional investors. This is reflected in the transparency of their business with their entire historical loan book available to download from their webpages.

Their website contains some informative examples that illustrate their appeal to customers. A particular strength seems to be the ease of use of their online systems. Another technological feature is that MarketInvoice, unlike factoring companies, does not take ownership of invoices or take responsibility for chasing repayment. Rather the security of their lending is cash flow based, by having the invoices paid directly into the same bank account where marketinvoice provides a line of credit. As invoices are paid down then the line of credit is automatically paid down with additional credit offered as new invoices are accepted by MarketInvoice.

- Funding Circle

Funding Circle is one the leading small business loan platforms, being praised as a '*real success story for British fintech*' by the UK Chancellor of the Exchequer, Philip Hammond (Turvill 2017). The platform was launched in 2010 (Collinson 2010) focusing exclusively on SMEs. Funding Circle has worked with a substantial number of prominent organisations. The UK Government started lending through the platform in 2013. In June, 2016, a £100 million lending partnership was agreed with European Investment Bank and a \$1 billion lending partnership with Alcentra in the US. The commercial arm of the government-owned British Business Bank in 2017 pledged to lend £40 million to UK SME through Funding Circle (British Business Bank 2017). In October 2018 it had an IPO on the London Stock Exchange.

It has provided \$5.8 billion in funds to 54,000 UK businesses (Funding Circle 2019a). However, Funding Circle's biggest market is the UK, however, the platform has extended its presence globally. Their 2018 annual report states that as of 31 December 2018, cumulative lending (from 2010) was \$1.9 billion to US borrowers, £ 4.6 billion to UK borrowers, €210 million to German borrowers and €140 million to borrowers from the Netherlands. The 2018 annual report goes on to suggest that it has £3.1 billion

loans outstanding as of 31 December 2018. This represents a 49% growth from 2017 (Funding Circle 2019b). Loan originations grew 40% to £2.3 billion in 2018. As of 31 December 2018, Funding Circle's typical customer across all countries has 10-year trading history, £950,000 of annual revenue and 8 employees. The average loan size and loan terms were £70,000 and 52 months, respectively (the loans range from 6 to 60 months). 11% was the average interest rate on loans originated from their platform (Funding Circle 2019a).

4.3.1 Trade Finance

Trade finance is another area that technology is supporting new finance. Trade finance is one of the most fragmented areas in the world of finance. Trade finance involves multiple documents and multiple parties which are largely paper-based and require substantial manual processing. Furthermore, complexity increases due to substantial regulatory compliance.

Kavuri and Milne (2019) find that trade finance is one of the most promising areas for technology-based provision of working capital. Costs are significant and technologies such as distributed ledger technologies have the potential to address the numerous inefficiencies arising from the failure to share data. This has led to various competing initiatives in the area which include:

- HSBC's VoltraOne platform (now being developed to a commercial scale by CryptoBLK). This aims to automate the entire trade finance process, and purely digital versions of the principal trade finance instruments.
- TIX platform of TradeIX and Marco Polo. It aims to improve transparency in the bank accounting systems that track customer trade finance.
- IBM supported we.trade which is based on the Hyperledger Fabric. It focuses on helping small and medium-sized enterprises access financing solutions.
- The Hong Kong Trade Finance Platform established as a co-operation between seven Hong Kong banks in 2017 in an effort to digitise all aspects of trade finance.
- Komgo initiative developed for the specialised area of commodity finance (oil, metals, agricultural commodities). They have the advantage of working in a specialised area with small number of large commodity traders and their banks, making the problem of co-ordination on a single solution relatively easy.

Despite this promise and several competing data sharing solutions, there is as yet no widespread adoption of technology-based international trade finance. A concern is that much international trade is transactional, with no guarantee that the same parties (importer, exporter, providers of finance) will continue working together. The costs of investing in access to a technology platform are therefore comparatively high compared to the gains from current client relationships and likely to have only limited payback. It will require much greater interoperability between all the different platforms for trade supply chains and trade finance in order to persuade the many parties involved to invest in the adoption of new processes. Thus, greater standardisation seems to be key to supporting full adoption.

5 Conclusions

Recent years have seen an explosion of new forms of technology-based credit. In this chapter we have focused on the provision of technology-enabled working capital finance to SMEs, looking at the overall magnitude of this lending and exploring the range of different ways in which this is being provided.

What emerges from our analysis is a picture of great variety of new products and business, but also that these are as yet only providing a relatively small proportion of the overall supply of working capital finance. Technology-based credit platforms provide around 2% of the stock of SME bank lending in the UK, a slightly higher proportion in the US and much less in almost all other countries. China, as so often in financial technology, has some of the most adventurous business models for working capital finance but again still small compared to conventional bank lending.

The new technologies are proving relatively successful at providing short-term credit to the very smallest 'micro-enterprises' in China, the US and the UK. Companies such as OnDeck, Kabbage and Funding Circle and others are successfully meeting demand for working capital finance from firms that previously would have been too small to access the same credit from conventional banks. An indication of the high levels of untapped demand are the high-interest rates and hence substantial margins achieved from this lending (albeit with some uncertainty about the underlying risks since these new business models have never been tested in a recession).

These new lenders are also pioneering much greater convenience and improved customer experience in loan applications, ideal for the smallest businesses that are under individual control and do not have the detailed

accounting and business modelling that is expected in application for conventional bank business loans. We can anticipate, over time, similar developments in other countries and, also, that banks will respond, either by acquiring loans from platform-based lenders or acquiring or developing their own more convenient customer interfaces (the cloud-based Biz2X software offering from Biz2Credit is an example of the potential for shared processing).

We can compete our analysis with a brief and preliminary policy discussion. Can technology ensure a much greater supply working capital: offsetting the contraction that has followed the post-crisis re-regulation of banking and avoiding the possibility of another collapse in the supply of working capital finance such as that which contributed so substantially to the post-crisis global downturn of 2018–2019? Can this increase the supply of the more sophisticated forms of working capital finance set out in Table 2, providing credit for longer periods or at lower risk and cost than the current generation of technology-enabled short-term high-margin lines of credit being provided to microbusinesses?

Our judgement is that this can happen, but it will require vigorous and far-sighted policy intervention, in order to support a much greater range of alternatives to bank provision of working capital finance. We can refer to Open Banking in the UK as an example of the kind of measures that are likely to be necessary. Open Banking means that bank customers are able to securely release their banking data to third parties, including non-bank credit providers, opening up a range of new potential customer services.¹⁹ This has employed competition law to force UK banks to work together to produce standardised application programming interfaces (APIs) that will work across the industry to support Open Banking. A fully competitive market in working capital finance will in our view require similar technological standardisation to support open access to and analysis of accounting records and for the digitisation and e-documentation of purchase orders and invoices.

The struggle of Amazon Lending to gain traction indicates that a pure ‘big data’ strategy based on internet traffic is not enough to offset the advantages that banks have in providing working capital finance through their access to banking transaction information. Without a broad policy to promote standardised access to accounting and supply chain as well as banking data, non-bank lenders will continue to face an un-level playing field against banks in most aspects of working capital finance.

¹⁹ For discussion of the potential benefits of UK Open Banking, see The Open Data Institute and Fingleton (2019).

Bibliography

- Bary, Emily. 2019. Fintech Unicorn Kabbage ends Q3 with nearly \$100 million in revenue amid loan growth. *Marketwatch*, October.
- British Business Bank. 2017. British Business Bank expands partnership with funding circle. *British Business Bank*, January.
- Business Expert. 2019. What is invoice financing and how can it benefit business cash flow? Business Expert.
- Chor, Davin, and Kalina Manova. 2012. Off the cliff and back? Credit conditions and international trade during the global financial crisis. *Journal of International Economics* 87 (1): 117–133. <https://doi.org/10.1016/j.jinteco.2011.04.001>.
- Claessens, Stijn, Jon Frost, Grant Turner, and Feng Zhu. 2018. Fintech credit markets around the world: Size, drivers and policy issues.
- Collinson, Patrick. 2010. Peer-to-peer lending and saving: Making everyone happy. *Guardian*, August.
- Dahl, Darren. 2015. The six-minute loan: How Kabbage is upending small business lending—And building a very big business. *Forbes*, May.
- Dean, James. 2017. Amazon Lending challenges banks by offering loans for small business. *The Times*, June.
- Federal Reserve Board. 2017. Report to the Congress on the availability of credit to small businesses. <https://www.federalreserve.gov/publications/files/sbfreport2017.pdf>.
- FSB. 2019a. Evaluation of the effects of financial regulatory reforms on small and medium-sized enterprise (SME) financing. Basel. <https://www.fsb.org/2019/11/evaluation-of-the-effects-of-financial-regulatory-reforms-on-small-and-medium-sized-enterprise-sme-financing-final-report/>.
- FSB. 2019b. Global monitoring report on non-bank financial intermediation 2018. Basel. <https://www.fsb.org/wp-content/uploads/P040219.pdf>.
- Funding Circle. 2019a. Funding Circle.
- Funding Circle. 2019b. Funding Circle 2018 annual report.
- Gawrych, Chloe. 2019. 11 best small business loans of 2020. Business.Org. <https://www.business.org/finance/loans/best-small-business-loans/>.
- Kabbage. 2019a. Where small business gets funded.
- Kabbage. 2019b. Kabbage launches Kabbage Payments™ to help small businesses get paid faster. October.
- Kavuri, Anil Savio, and Alistair Milne. 2019. Evolution not revolution: An assessment of crypto assets and distributed ledger technologies in the banking sector: Short version.
- Market Finance. 2019. What is working capital—a guide to working capital solutions.
- OnDeck. 2019a. Because getting a business loan shouldn't feel like a second job.
- OnDeck. 2019b. Investor relations. Ondeck.
- OnDeck. 2019c. Ondeck 2018 annual report.

- OnDeck Capital. 2018. OnDeck small business online lending tops \$10 billion. *Prnewswire*, September.
- Perez, Sarah. 2019. This Week in Apps: Apple Arcade's new franchise, Fortnite takes on Google Play, the Disney+ app footprint. *Techcrunch*, December.
- Prosser, Marc. 2015. The difference between purchase order financing and invoice factoring. *Bizjournal*, September.
- Rhodes, Chris. 2018. Business statistics. 06152. Briefing Paper. London. <https://researchbriefings.files.parliament.uk/documents/SN06152/SN06152.pdf>.
- The Open Data Institute, and Fingleton. 2019. Open Banking, preparing for lift off. London. <https://www.openbanking.org.uk/wp-content/uploads/open-banking-report-150719.pdf>.
- Turvill, William Turvill. 2017. Chancellor Philip Hammond praises 'success story' funding circle as it announces \$100m investment. *City A.M.*, January.
- Zhang, Bryan, Tania Ziegler, Leyla Mammadova, Daniel Johanson, Mia Gray, and Nikos Yerolemou. 2018. The 5th UK alternative finance industry report. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2018-5th-uk-alternative-finance-industry-report.pdf.
- Ziegler, Tania, Daniel Johanson, Michael R. King, Bryan Zhang, Leyla Mammadova, Feliipe Ferri, Roswitha Trappe, et al. 2018. Reaching new heights: The 3rd Americas alternative finance industry report. Cambridge. https://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2019-05-ccaf-3rd-americas-alternative-finance-industry-report.pdf.



Robo-Advising

Francesco D'Acunto and Alberto G. Rossi

Individual investors are known to make significant mistakes relative to the optimal behavior of a standard economic agent. Mistakes range from saving too little to maintain one's living standards after retirement to incorrect exposure to debt and equity instruments throughout one's life cycle. Even though financial advisers have traditionally been considered the main solution to limit the impact of such mistakes on financial decisions, a set of limitations of human advisers have also been documented, ranging from the transmission of their own personal biases into investors' portfolio choices to the high costs of financial advice, which make this form of advice not accessible to large fractions of consumers/investors. Over the last decade, robo-advising—automated algorithmic financial advice—has emerged as a potential solution to these limitations.

In this chapter, we first discuss the limitations of traditional financial advice, which led to the emergence of robo-advising. We then describe the main features of robo-advising and propose a taxonomy of robo-advisors

F. D'Acunto (✉)

Carroll School of Management, Boston College, Boston, MA, USA
e-mail: dacunto@bc.edu

A. G. Rossi

McDonough School of Business, Georgetown University, Washington, DC, USA
e-mail: agr60@georgetown.edu

based on four defining dimensions—personalization, discretion, involvement, and human interaction (Sect. 1). Building on these premises, we delve into the theoretical and empirical evidence on the design and effects of robo-advisors on two major sets of financial decisions, that is, investment choices (for both short- or long-term horizons) and the allocation of financial resources between spending and saving (Sects. 2 and 3). We conclude by elaborating on five broadly open issues in robo-advising, which beget theoretical and empirical research by scholars in economics, finance, psychology, law, philosophy, as well as regulators and industry practitioners (Sect. 4).

1 From Traditional Financial Advisers to Robo-Advisors

1.1 Limitations of Traditional Financial Advisers

In principle, financial advisers could help investors make better investment choices. Delegating individual portfolio allocations to an adviser who manages multiple portfolios allows for economies of scale, whereby the costs of information acquisition advisers incur are shared across all clients. Also, advisers' financial literacy and skills might be higher than those of most individual investors.

In practice, though, a variety of conflicts of interest plague the client–adviser relation. For instance, Hackethal et al. (2012) show advised accounts perform worse than unadvised accounts largely because advised accounts trade more often, producing commissions for advisers at the expense of investors. Moreover, Linnainmaa et al. (2021) show that advisers make the same mistakes in their own investment accounts as they do in their clients' accounts, and hence transmit their own biases to clients. These results cast doubts on whether traditional financial advisers can create value for the investors whose accounts they manage.

At the same time, traditional financial advisers might be beneficial to individual investors in ways that differ from the implementation of more profitable investment strategies. For instance, Gennaioli et al. (2015) propose a model whereby investing through financial advisers may be rational because it increases risk-taking by generating “peace of mind” in investors. Reducing risk-aversion in financial choices might increase investors' returns enough to compensate for the fees advisers charge. Foerster et al. (2017) confirm empirically that higher trust in advisers results in higher risk-taking by investors. Higher experienced returns, though, are not enough to compensate for the

higher fees. Either investors do not know how much they pay for advice or they value aspects other than portfolio-return maximization when interacting with financial advisers. Consistent with the second possibility, after conducting a comprehensive survey to elicit clients' needs in financial advice, Rossi and Utkus (2019b) find that individual investors hire financial advisers largely to satisfy needs other than portfolio-return maximization. Such needs include acquiring "peace of mind," having access to the opinions of an expert, and delegating financial decisions. Also, most investors do not know how much they pay for financial advice.

1.2 The Emergence of Robo-Advising

Robo-advising has been emerging over the last two decades as an alternative to traditional financial advisers and as a way to address their limitations. Robo-advisors are digital platforms that provide financial advice to investors in an automated fashion. Depending on their degree of sophistication, robo-advisors collect and use individual-specific information to construct tailored financial plans and advice for investors, as we discuss in detail below.

Robo-advisors have the potential to benefit end-consumers for a number of reasons. First, they can offer financial advice against low fees, because investors' portfolio allocations are fully automated. Second, robo-advisors can serve individuals with any level of wealth, whereas human financial advisers are time-constrained and hence typically cater to wealthier households. Third, robo-advisors are based on automated algorithms that can be monitored and improved over time. Fourth, unlike the decisions of human advisers, the decisions of robo-advisors are hardwired and can be reviewed and explained to investors as well as to regulators consistently. In Sects. 2 and 3, we discuss the extent to which existing empirical evidence supports these potential benefits of robo-advising.

1.3 A Taxonomy of Robo-Advisors: Four Defining Features

The blanket term "robo-advising" hides a variety of different models and methods that differ among several dimensions, four of which are defining features:

1. **Personalization** of the advice;
2. **Involvement** of the investor in financial plans and choices;

3. Investors' **discretion** to deviate from the automated advice;
4. The presence of any form of **human interaction**.

We delve into each one of these aspects below.

1.3.1 Personalization of Advice

Human advisers claim they can tailor investment strategies to the individual needs of each investor, although mounting evidence casts doubt on this claim (Linnainmaa et al. 2021). Robo-advisors, instead, vary dramatically on the extent to which they can create individually designed investment portfolios and financial plans. Target Date Funds (TDFs) are one end of the spectrum. TDFs can be considered the first and most primitive form of robo-advising, in that the investment strategy they implement abstracts from all investor characteristics, except for investors' age. Depending on the year when investors plan to retire, investors purchase a cohort-specific TDF that is rebalanced automatically over time. The main strategic asset allocation TDFs perform is to reduce the wealth invested in equities and increase the wealth invested in fixed-income securities over one's life cycle.

More recent robo-advisors elicit qualitatively (or quantitatively) a set of demographic characteristics such as investors' income bracket, investment horizon, willingness to take financial risk, and job security. Robo-advisors then propose investment plans and strategies that are the same for each individual investor who falls in the same categorization based on these demographics. This level of personalization is common in US commercial robo-advisors such as Wealthfront, Betterment, and Vanguard's Personal Advisor Services (PAS). Whereas the multi-dimensional characterization of investors allows for more tailored strategies than TDFs, several important aspects that should determine an investors' strategy, such as non-financial investments or upcoming expenses such as children's college education are typically disregarded. Moreover, as we discuss in more detail in Sect. 4, the extent to which information can be elicited directly from investors, who tend to lack financial literacy, without a qualitative human assessment and only based on pre-designed bucketing is an open question in robo-advising. On top of demographic characteristics and risk preferences, robo-advisors designed for investors' trading in individual stocks and short-term investing generally add individuals' existing portfolio allocations to the inputs used to generate optimal weights for portfolio allocations.

In the context of personalization, the trade-off of robo-advising rests between providing tailor-made solutions that are specific to each investor, but

could potentially result in poor ex-post performance for some investors, and providing less extreme positions and thus more robust portfolio allocations, which though might fail to consider important unique investors' features at sign up and over time.

1.3.2 Investor Involvement

A second aspect that distinguishes different types of robo-advisors is investor involvement. Robo-advisors for trading usually ask investors to approve every single trading decision before it is executed. In this way, investors can modify the course of action the algorithm suggests and require a re-optimization of their financial plan and strategy at any point in time. D'Acunto et al. (2019c) analyze one such robo-advisor. This form of advice, in which investors are directly involved in approving or denying investment choices, should be labeled "robo-advising" in the strict sense. Indeed, these robo-advisors provide algorithmic advice to the investor and make the implementation of advice extremely simple, for instance by producing automatically all the trades needed to implement a portfolio rebalancing strategy and allowing investors to implement the advice by simply clicking on a button (D'Acunto et al. 2019c). Ultimately, though, any decision-making authority rests with the investor.

At the other end of the spectrum are robo-advisors for long-term investing, who not only provide automatically generated financial plans and strategies, but also place trades automatically on behalf of investors. These robo-advisors request the approval of an initial plan, but once the investment plan has been approved, they manage investors' wealth and trade without any input from investors. A more correct taxonomy would define such form of automated advice "robo-managers" rather than robo-advisors, because robo-managers manage investors' wealth directly rather than providing advice about each step in the implementation of the strategy.

In terms of investor involvement, the main trade-off in robo-advising consists in either allowing investors to retain full control of their portfolio at the expense of paying attention to the management of their wealth, or replacing fully the individual as a decision-maker with an algorithm.

1.3.3 Investor Discretion

Discretion is investors' ability to override robo-advisors' recommendation. Robo-advisors that allow for more discretion let investors modify the portfolio weights the algorithm proposes. In other cases, investors can also choose whether the trades proposed by the algorithm should be implemented. Finally, investors might include stocks and other assets in the financial plan not recommended by the robo-advisor, in which case the robo-advisor would optimize the portfolio formed of the mix of recommended assets and investor-proposed assets (D'Acunto et al. 2019c).

Other robo-advisors allow portfolio weights personalization only within certain pre-set guardrails (see Rossi and Utkus 2019a). They allow investors to take more or less risk, relative to what the algorithm suggests, but are inflexible in terms of picking what parts of the plan to accept and what part of the plan to reject. For instance, fully automated robo-advisors such as Wealthfront and Betterment allow little discretion.¹ Hybrid cases include Vanguard PAS allow investors to voice their preferences with a human financial adviser who, in turn, has the power to override some of the allocations suggested by the algorithm manually.

Client discretion constitutes an important trade-off for robo-advising design. Robo-advising is a paradigm that lies in between pure *libertarianism*, in which individuals are left on their own to make investment decisions, and *libertarian paternalism*, in which individuals are defaulted into what economists believe is the best decisions for them based on standard economic theory, and individuals can only choose to opt out of the assigned defaults. The more discretion is programmed into the robo-advisor, the more libertarian the robo-advisor is. At the extreme where the individual is not granted any discretion, robo-advising can be thought as a form of libertarian paternalism, where investors give up the possibility to make their own individual decisions in managing their wealth and comply to an optimized option chosen based on the prescriptions of standard economic theory.

1.3.4 Human Interaction

The last differentiating feature among robo-advisors is the degree of interaction investors have with human advisers, if any. Many robo-advisors, such as Wealthfront and Betterment, are purely automated and investors cannot

¹ In fact, Betterment recently launched a new product named Flexible Portfolios to cater to the investor that wanted additional discretion in the management of their wealth.

access any human advisers. Not employing human advisers allows these robo-advisors to maintain their operating costs low. Fully automated robo-advisors cater to younger cohorts such as millennials, who are generally comfortable with having their wealth managed by algorithms, without the presence of a human who explains the intricacies of how their wealth is managed.

Other robo-advisors that cater to a wealthier and older clientele are hybrid in nature. The majority of the heavy-lifting in terms of designing the portfolio allocation is performed by the algorithm, but human advisers interact with the investor at key moments, such as at sign-up as well as when investors have important questions about their portfolios. Human advisers' presence is crucial to ensure customers' needs are being satisfied as well as to handle all financial planning tasks—such as opening college funds or IRA accounts, for example—that are not easily automated. Human advisers are also crucial to keep investors participating in equities in periods of bear markets, when investors may become more fearful and may be trying to reduce their exposure to risk.

1.4 What Financial Decisions Do Robo-Advisors Aim to Improve?

1.4.1 Investment Decisions

The first issue robo-advisors aim to target is limited exposure to risky assets. Using the 2001 Survey of Consumer Finances, Campbell (2006) shows a low fraction of individuals at the lower end of the wealth distribution invest in public equity, even though finance theory predicts everybody—absent costs of participation—should participate in the stock market. Surprisingly, stock market participation is not widespread at the higher-end of the wealth distribution either. For example, at the 80th percentile, 20% of consumers are not exposed to the stock market.

A second major shortcoming of individuals' investment decisions relates to portfolio allocations and especially the pervasive lack of diversification. A consistent finding in the literature is investors tend to own few risky assets, which are typically individual stocks or specialized mutual funds. Barber and Odean (2001) report the median client in a large US brokerage house held only 3 stocks in the years 1991–1996. Using a US sample between 2013 and 2015, Gargano and Rossi (2018) find similar patterns—the median US investor held only 4 stocks. These findings are also true internationally—D'Acunto et al. (2019c) find the median investor in a large Indian brokerage house holds 5 stocks. Lack of diversification also arises in terms of lack of

exposure to varied geographic shocks. Individuals display local bias in their investments (e.g., see French and Poterba 1991). In particular, US investors have very little exposure to international markets and prefer to purchase local companies rather than companies headquartered in other states. US investors also hold a disproportionate amount of their employers' stocks within their 401(k) plans (Mitchell and Utkus 2004).

Active individual investors are also known to be subject to a variety of behavioral biases when they trade stocks. For instance, Odean (1999) and Barber and Odean (2001) argue investors trade too much. Odean (1998) shows investors are more likely to realize trading gains as opposed to trading losses (disposition effect). This behavior is sub-optimal, because selling gains are taxable items not netted by trading losses.

In Sect. 2, we discuss the extent to which robo-advising has been able to tackle these issues with individual investment decisions thus far.

1.4.2 Consumption-Saving Decisions

Possibly, the most important financial decision individuals make is how much to save for retirement. This is an extremely complicated choice that depends on expected future income growth, the probability of incurring unemployment spells, the expected equity premium, the sustainability of social security benefits, projected healthcare costs, and individual risk-aversion, among other characteristics of which consumers might not even be aware. Indeed, Lusardi and Mitchell (2007) find that consumers who answer basic financial literacy questions incorrectly are less likely to save enough for retirement. A study by Gomes et al. (2020) uses information on contribution rates, salary, investment plan features, and asset allocation for more than 300K individuals. They show three quarters of the workers in their sample are unlikely to be able to maintain their pre-retirement consumption after retirement.

Robo-advisors aim to tackle these issues in individual consumption-saving decisions by reducing consumers' lack of information about their own inflows and outflows as well as by nudging consumers' choices through individualized messages and other vivid nudges to obtain behavioral reactions that would be hard to implement in the absence of an online platform or internet application.

In Sect. 3, we discuss the applications of robo-advising to this realm and the empirical evidence on the effects of such applications.

2 Robo-Advising in Asset Management

The most developed applications of business-to-consumer robo-advising are in the realm of asset management. Robo-advisors have developed to target both short-/medium-term investment (“robo-advisors for trading”) as well as long-term investment and especially the phase of accumulation of resources for retirement (“robo-advisors for passive investors”).

Building on the characteristics of robo-advisors we described in Sect. 1, the crucial difference between robo-advisors for trading and robo-advisors for passive investors is that the former have higher personalization of investment advice, promote direct involvement of investors in the definition and implementation of strategies, and allow for discretion in terms of deviating from the advice. Active involvement requires robo-advisors for trading to provide vividly the information investors need to understand and process investment strategies, and thus provide also an educational role that is less relevant in robo-advisors for passive investors.

The second main difference between the two forms of robo-advisors is their focus on different asset classes. Robo-advisors for trading focus on individual stocks or highly specialized mutual funds, whereas robo-advisors for passive investors typically target exchange-traded funds (ETFs) and low-fee mutual funds. D’Acunto et al. (2019c), Reher and Sun (2019), and Rossi and Utkus (2019a) are among the first academic studies of the characteristics, mechanics, as well as direct and indirect effects of robo-advising for short-term and long-term investing. In the rest of this section, we discuss the main features of the design of robo-advisors in asset management and their effects on investment performance.

2.1 Robo-Advising for Short- and Medium-Term Investing

Robo-advisors for trading are based on Markowitz mean–variance optimization and aim to maximize portfolios’ Sharpe ratios. D’Acunto et al. (2019c) study a *Portfolio Optimizer* targeting Indian equities, which has a similar scope as US robo-advisors for trading such as *MI Finance*. In terms of design, the optimizer displays three main features: (i) the feeding of information about the expected returns of individual securities developed by the brokerage house’s research team; (ii) the estimate of the variance–covariance matrix based on three years of historical daily observations; and, (iii) the use of shrinkage techniques and short-sale constraints to limit estimation-error effects and guarantee well-behaved portfolio weights.

The robo-advisor produces automatically the buy and sell trades the investor would need to perform to implement the advice, and the investor can place the trades automatically in batch mode by simply clicking on a button on the application's screen. The simplicity of execution of the advice is a fundamental feature that distinguishes robo-advising for trading from other forms of electronic investment advice, which require investors to come up with the implementation of the advice in their own portfolio and for this reason are often ineffective (Bhattacharya et al. 2014). This robo-advisor for trading provides substantial personalization of advice, which is partly based on investors' own starting portfolio of assets and allows for discretion in the assets users want to incorporate in their portfolios.

Using viable counterfactuals in a difference-in-differences design, D'Acunto et al. (2019c) find that robo-advice for trading is beneficial to ex-ante undiversified investors, because it increases the diversification of their portfolios hence reducing portfolio volatility. It also produces slightly higher ex-post mean returns. At the same time, the robo-advisor does not improve the performance or volatility of the portfolios of already-diversified investors. If anything, due to the higher amount of trading when rebalancing their portfolios and the frequency with which active investors engage in rebalancing, average after-fee returns are lower for diversified investors after they start to use the portfolio optimizer relative to before.

A crucial dimension under which robo-advisors differ from human financial advisers is the extent to which advisers' biases and misguided beliefs can be transmitted to investors' portfolios. The robo-advisor judges all potential trades based on the underlying algorithm. Behavioral biases common among individual investors should have little scope under robo-advising, as long as the developers of the algorithms did not embed such biases into their codes. And, indeed, D'Acunto et al. (2019c) find that the incidence of well-known biases, such as the disposition effect, the rank effect, and trend chasing, decreases for all investors after accessing robo-advising, irrespective of their characteristics and levels of diversification.

2.2 Robo-Advising for Long-Term Investing and Retirement

Robo-advisors for long-term investing target mostly indexed mutual funds and ETFs. The main objective of these robo-advisors is to set risk-factor exposures that are compatible with investors' preferences and investment horizons. These robo-advisors do not engage in stock picking, but trade to rebalance portfolios at pre-set regular intervals (usually quarterly). These robo-advisors

usually move investors in and out of risky assets as a function of one's time to retirement and do not engage in market timing by increasing equity exposures when expected returns are high and reducing equity exposures when expected returns are low.

A defining feature of robo-advisors for long-term asset management is that they not only provide advice on rebalancing strategies, but they directly manage investors' portfolios. They require minimal to no involvement from investors. In fact, these robo-advisors emphasize that investors should only worry about contributing resources to the managed portfolios without even paying attention to their finances (Gargano and Rossi 2018).

2.2.1 Targeting Performance and Sharpe Ratios

Rossi and Utkus (2019a) study the effects of a large U.S. hybrid robo-advisor on the portfolios of previously self-directed investors. They find that, across all investors, robo-advising reduces holdings in money market mutual funds and increases bond holdings. It also reduces idiosyncratic risk by lowering the holdings of individual stocks and US active mutual funds and raising exposure to low-cost indexed mutual funds. It further eliminates home bias by significantly increasing international equity and fixed-income diversification. Finally, over the sample period the authors analyze, robo-advising increases overall risk-adjusted performance, which is largely driven by lower portfolio risk.

Rossi and Utkus (2019a) use a machine-learning algorithm, known as Boosted Regression Trees (BRT), to explain the cross-sectional variation in the effects of robo-advising on portfolio allocations and performance. Investors who benefit from advice are those with little self-directed investment experience at managing their wealth, with large cash holdings, and with high trading volume before adopting robo-advising. Investors with little mutual fund holdings and investors invested in high-fee active mutual funds also display significant performance gains. Moreover, investors who end up benefiting more from robo-advising are more likely to sign-up and less likely to quit the service over time.

Reher and Sun (2019) study a large US-based robo-advisor for long-term investing, which aims to optimize investors' portfolios Sharpe ratios. They find that underdiversification increases the likelihood of uptaking robo-advising as well as the amount of deposit inflows, especially for middle-class investors, thanks to the low minimum account sizes. The typical drop in minimum account size for the robo-advisor relative to the standard investment accounts, which amounts to a 90% drop in minimum account size,

leads to a 56% increase in account flows—new accounts opened by less wealthy individuals. In terms of average performance, users' Sharpe ratios increase by about 10% after the uptake of robo-advising, and the bulk of the improvement in performance is driven by a drop in portfolio's exposure to idiosyncratic risk and a sharp reduction of volatility. Overall, the benefits of diversification and the less stringent requirements to access robo-advising relative to traditional financial advice are the crucial drivers of investors' gains from robo-advising. Conversely, these results, similar to D'Acunto et al. (2019c) for the realm of short-term investing, might suggest that the investors who are most likely to gain from adopting robo-advising are the less financially literate investors.

Several robo-advisors propose specialized features that might attract wealthier and more financially sophisticated investors. For instance, they embed functions to optimize tax-loss harvesting—the practice of replacing an asset at loss within the portfolio with similar assets in terms of expected return and volatility so as to offset capital gains and reduce tax debentures. By construction, tax-loss harvesting is only relevant to investors who hold taxable brokerage accounts, and hence such functions do not target retirement accounts that fall under the types that allow for deferred taxation. Moreover, tax-loss harvesting is obviously only relevant to investors who face high marginal tax rates and for whom offsetting capital gains and capital losses might produce nonnegligible benefits in terms of reduced tax debentures. Because tax-loss harvesting requires monitoring one's portfolio and outside opportunities often throughout the fiscal year, this feature is marketed as a benefit of robo-advising for wealthy investors, who might have a higher opportunity cost of time. In the US, both *Betterment* and *Wealthfront* emphasize their tax-loss harvesting focus. For example, Betterment's "Tax Loss Harvesting+" algorithm checks daily for harvesting opportunities and trades stocks. The algorithm also reinvest every harvested dollar in asset classes that bring the client portfolio back into balance rather than defaulting back to the original asset class.

2.2.2 Targeting Risk Levels: Value-at-Risk (VaR) Strategies

The robo-advisors discussed above propose simple techniques based on Markowitz's principles of mean–variance optimization and in which the assessment of expected returns for individual stocks or funds is a crucial input of investment strategies. Recently, robo-advisors with additional features have been developed. For instance, *scalable.CAPITAL*, a robo-advisor founded in Germany and diffused across several European countries, aims to provide

active risk management to portfolios that mainly invest in passive asset classes such as ETFs and mutual funds. The robo-advisor uses a VaR approach in their risk-management strategy, and is targeted to investors whose level of financial literacy would not allow understanding and/or designing an active risk-management strategy. VaR strategies consist of fixing a maximum allowed yearly percentage loss of portfolio value with a high confidence level, usually 95% percent. The robo-advisor models dynamically the expected probability of losses and adjusts the portfolio composition automatically to avoid deviations from the pre-set loss thresholds.

The robo-advisor allows discretion by enabling users to choose the VaR thresholds and other features over time, based on changing circumstances in investors' financial conditions. Moreover, human interaction is allowed in two ways—customer service teams who are specialized in answering questions and concerns about investment strategies and execution, as well as a human team who oversees the execution of trades.

Finally, *scalable.CAPITAL* also includes educational components to enhance investors' understanding of the VaR strategy the robo-advisor uses and their understanding of general principles of financial econometrics. This educational function is implemented through simple and vivid descriptions of the concepts of confidence intervals, portfolio loss, and VaR at the time in which users decide the levels of risk they want to face in their portfolio, as well as through the availability of podcasts that explain the basic statistical concepts behind the robo-advisor's strategies, which can be accessed online and from investors' applications.

3 Robo-Advising and the Consumption-Saving Decision

Whereas the most common applications of robo-advising in the recent years aim to help households in the realm of asset management, more recent applications have targeted the daily consumption/saving decisions consumers make, which represent the foundation of any life-cycle model of allocation of financial resources over time. Indeed, the problem of allocating resources between alternative forms of investment, whether in the long or short term, only arises after the allocation between spending and saving is set.

The applications of robo-advising to daily spending and saving choices have developed in response to well-known departures from the standard problem of consumption allocation throughout the life cycle. Such departures have been documented in the field and the laboratory by a large literature

in behavioral economics and social psychology. After all, solving an optimal allocation problem throughout one's life cycle to determine the share of spending over income at each point in time goes beyond the cognitive abilities and literacy skills of most consumers (D'Acunto et al. 2019b). For this reason, consumers use rules of thumb, which are often polluted by the lack of information and the formation of distorted beliefs about relevant economic variables, the prevalence of non-standard preferences and beliefs such as present bias and self-control, or the persistence of social norms of behavior to which households conform without critical assessment (for instance, see Guiso et al. 2008; D'Acunto et al. 2019d).

3.1 Robo-Advisors Targeting Consumers' Informational Frictions

The first type of demand-side frictions robo-advisors have targeted is informational in nature. Most households do not collect regularly information about the economic variables that should drive their choices, or are unable to understand and elaborate this information even when it is provided to them (D'Acunto et al. 2019a). Robo-advisors have developed strategies to target this issue directly by making it easy and practical for consumers to form beliefs about such important economic dimensions. The most basic informational friction robo-advisors have targeted is the lack of information about one's own balance sheet in terms of inflows and outflows. Similar to corporation, most households face a mismatch in the liquidity of inflows and outflows (e.g., monthly salaries against one-time large expenditures of durable goods) or in the economic horizon of investment decisions (e.g., choice of spending on durable goods at the time of purchase or with delayed instalments). These mismatches are exacerbated by the fact that several households have their inflows and outflows split into several separate accounts, such as checking accounts in which their inflows are transferred and one or more credit card accounts that collect the outflows.

To alleviate this complex budgeting problem, robo-advisors have developed in the form of *income aggregators*. Income aggregators are desktop or phone applications, in which users link their financial accounts, including checking, credit card, and investment accounts. The main function of income aggregators is to elaborate all the transaction-level information they obtain from the individual accounts in one single organized balance sheet, which provides households with a clear and immediate overview of their own financial situation at each point in time (Baker 2018).

On top of this aggregation function, which by itself reduces consumers' lack of information about their own finances, several robo-advisors have designed specific information interventions that warn users about abnormal inflow or spending patterns through the use of notifications. Lee (2019) studies one of such interventions—a FinTech application that provides users with notifications every time their spending patterns increase abnormally relative to average spending. In this way, users face an immediate and vivid nudge to adjust their spending to their own average level, and indeed Lee (2019) finds that users respond to such nudges.

Providing information about one's own finances might not be enough to correct potential distortions in consumption and saving choices relative to the choices of neoclassical economic agents. This issue arises because often households use rules of thumb to assess the ratio of spending over saving in their daily life. A common rule of thumb is the conformity to peers' spending and saving. That is, consumers might think that the observed patterns of spending by peers contain information about their own optimal pattern of spending. Distortions in choice might arise especially in times of social media, in which the most conspicuous part of peers' consumption is public thus causing a *visibility bias*, which makes agents believe that their peers spend more, on average, than what they really spend (Han et al. 2019).

Recent applications of robo-advising have tackled this informational friction. For instance, *Status Money* is an application that, on top of the baseline income aggregation features, provides each user with information about the average spending, assets, debts, and net worth of individuals that have similar demographic characteristics, whose information is crowdsourced using transaction-level data from a large, representative US population. D'Acunto et al. (2019f) study the effects of this intervention on spending behavior. They observe the spending of users in the months before and after sign up and find that on average users converge to the spending level of peers disclosed by the app, thus indicating that users find the signal they obtain about peers informative. Interestingly, users that appear to spend more than their peers react systematically more than users that appear to spend less than the peers—in the overspending domain, users on average reduce their spending by 3 percentage points of their monthly income, which amounts to about \$247, whereas those in the underspending domain increase their spending by 1 percentage point of income on average. The authors find that the information content of the signal is important to trigger households' reaction, because the most reactive users are those for which the peer group is based on more similar demographic characteristics, relative to those who face peer groups defined on broader demographic characteristics.

3.2 Robo-Advisors Targeting Consumers' Non-standard Preferences and Beliefs

The second type of demand-side friction derives from the cross-sectional diffusion of preferences and beliefs that do not adhere to the standard neoclassical framework. A common example is present bias deriving from hyperbolic discounting: if economic agents discount consumption in the future more than present-day consumption, they will tend to overconsume at each point in time at the expense of saving and hence future spending (Laibson 1997).

Several robo-advising applications have developed to help agents correct present bias across a large set of domains, such as the provision of electronic messages on the balance of Supplemental Nutrition Assistant Programs (SNAP) for low-income households (Hillis 2017), the use of AI algorithms to predict which consumers might incur in future overdraft fees and warning such users with carefully framed messages (Ben-David et al. 2019), or the use of different framings and designs of information provision to nudge households' spending and saving behavior (Levi 2019; D'Acunto et al. 2019e; Gargano and Rossi 2019).

Recent studies have used the laboratory and synthetic markets to test alternative choice architectures and framings for robo-advisors targeting choice inertia in the realm of financial planning. For instance, Jung and Weinhardt (2018) find that defaults and warning messages reduce financial decision inertia, and uncover interesting differences in inertia in financial decision-making across genders.

4 Open Issues in Robo-Advising

So far, we have focused on a positive analysis of the features and characteristics of robo-advising services. The unprecedented and swift evolution of financial advice propelled by algorithmic applications, though, also proposes a set of issues that are still broadly open questions for researchers, practitioners, and regulators alike. These broad and interdisciplinary questions require attention by scholars in finance, economics, law, social psychology, and philosophy. In this section, we give an overview of the open questions and propose directions for future research.

4.1 From Domain-Specific to Holistic Robo-Advising: Across Realms and Over the Life Cycle

Since its origin, theoretical research on optimal consumption, spending, and investment choices has emphasized the holistic nature of the optimal life-cycle allocation of resources (e.g., see Carroll 1997, 2000). This holistic allocation encompasses two dimensions—(i) the optimal allocation of resources across different realms at each point in time, such as the share of wealth allocated to pay mortgage loans, student loans, and invest in retirement savings during one's working life, as well as (ii) the optimal allocation of resources throughout one's life cycle, from the early stages of investing in human capital and education to the decumulation phase after retirement and the allocation of bequests.

Human advisers, despite the limitations we have discussed above, aim to propose such a holistic approach to financial advice. Instead, the majority of existing robo-advisors focus almost exclusively on one or a few limited domains. Robo-advisors targeting retirement investment, for instance, barely ever provide advice on mortgage uptake, student-loan assessments, or the timing and viability of large durable consumption spending. The design of a holistic robo-advising service requires research on both the theoretical and empirical side. On the theoretical side, existing models of optimal life-cycle consumption, saving, and investment decisions do not delve into all the peculiarities of spending opportunities, or the different types of investments (education, large durable goods, housing) and associated forms of financing (Browning and Crossley 2001). More realistic theoretical approaches have incorporated two or three features at once (e.g., see Cocco 2005; Cocco and Gomes 2005), and more progress is needed to guide empirical applications. In particular, the phase of decumulation after retirement has obtained little attention, which translates into robo-advisors that may fail to address the complex choices retirees have to make. Baker and Dellaert (2019) propose a framework to develop this underexplored dimension.

Designing a holistic robo-advisor also faces empirical challenges. Procedures that truly allow individual-specific tailoring of advice are still lacking. Most existing robo-advisors place users in broad categories related to the willingness to take risk, age profile, and a few other demographics. Users that fall into the same buckets obtain the same advice, although obviously each user might differ under important non-elicited dimensions or the buckets might be too broad to capture preferences and beliefs accurately. The first step is to understand whether the dimensionality of this problem can be reduced by determining which characteristics are more

important to be targeted and hence elicited by robo-advisors. Traditional empirical methods can barely help, e.g. linear regressions of investment outcomes on a kitchen-sink list of potential characteristics would face the issue of overfitting and ultimately be uninformative. Rossi and Utkus (2019a) make progress by using machine-learning techniques to assess which individual characteristics explain more of the variation in investment performance across investors, as well as what is the actual (non-linear) relationship between each characteristic and investment outcomes. Another approach is to ask consumers directly which features of advice are important to them (Rossi and Utkus 2019b). Future research should provide additional evidence on which characteristics robo-advisors should target, as well as new methods to elicit these characteristics even if many investors are financially illiterate and unable to express their own economic preferences and beliefs consistently (e.g., see D'Acunto et al. 2019a, c). For instance, Alsbah et al. (2019) propose a reinforcement learning framework in which the robo-advisor does not need to be fed rough qualitative risk preferences of clients manually, but learns risk preferences over time by observing portfolio choices under different market conditions.

The second empirical challenge is the design of data-analytic methods to analyze multi-faceted information encompassing several aspects of one's financial profile at once. Recent providers suggest potential solutions. For instance, *Pefin*, a US-based holistic robo-advisor, uses a feed-forward neural network whose input consists of a broad set of aggregate characteristics (e.g., macroeconomic variables, financial regulatory changes) and individual-user characteristics (e.g., changes in spending and saving profiles) that are allowed to change over time. In this way, *Pefin* proposes a continuously changing set of financial projections and updating financial plans, which instead are typically static for most robo-advisors. More research on the design of such applications and the performance of investors that follow such holistic advice relative to viable counterfactuals is imperative to make progress in this area.

4.2 Algorithmic Aversion: Is Hybrid Robo-Advising a Solution?

Do consumers trust advice coming from a machine, with which, contrary to human advisers, no empathic interactions are possible? Two points suggest that the distrust toward machines, a.k.a. algorithmic aversion, is likely an important issue that deserves further study in finance, social psychology, marketing, and related fields. First, research in social psychology on the extent of users' trust in algorithmic vs. human judgment delivers conflicting results

(Logg et al. 2019; Castelo et al. 2019). How robo-advisors could promote algorithmic appreciation instead of aversion is an important open question. Second, the fact that younger users are more likely to access digital advice than older users (e.g., see Sironi 2016; Ben-David and Sade 2018; D’Acunto et al. 2019f) suggests a strong divide between digital native users and others. This age profile might fade as the new generations grow, but a challenge for the next decades is to understand why middle-age and elderly consumers avoid robo-advice and which interventions might increase uptake. After all, older investors currently own the majority of wealth in the economy, whereas digital natives have barely yet started to accumulate any wealth for retirement.

An important role of human advisers that robo-advising can barely fulfill is that of money doctors, i.e. the responsabilization of the human adviser that leads to investors accepting higher risk in individual portfolios (Gennaioli et al. 2015; Rossi and Utkus 2019b). Hybrid robo-advisors—robo-advisors whose strategies and planning are fully automated, but allow users to interface with a human being—have been proposed as a solution. In hybrid robo-advising, the role of humans is only slightly more involved than in a customer desk of a service company. Several questions about hybrid robo-advising are still open. How do the uptake of robo-advising and hybrid robo-advising differ? Do demographics that trust algorithms less really increase uptake when robo-advisors are hybrid? To what extent does the hybrid nature also fulfill an educational role, whereby financially illiterate investors might learn and use such knowledge in other realms of economic decision-making? Can humans reduce the high drop-out rates of robo-advising users, especially in times of bear markets? Research using observational and experimental methods should inform the role of hybrid robo-advisors positively and normatively.

A specific form of hybrid robo-advising that has obtained interest in the industry is the *super adviser*. Super advisers are human financial advisers who make use of robo-advising to produce financial plans and strategies, but represent the *only* interface between users and their investment strategies and performance. Super advisers resemble traditional human advisers on the client side, but are closer to robo-advisors in charging lower fees—because super advisers do not need to spend time producing financial plans, strategies, and implementing such strategies—and better performance, on average. Super advisers could also represent a solution to the transmission of human advisers’ own biases and misguided beliefs to their clients’ portfolios (Linnainmaa et al. 2021). Understanding the pros and cons as well as the costs and benefits of super advisers requires research designs that provide viable counterfactuals.

4.3 Will Robots Democratize Access to Financial Advice or Exacerbate Inequalities?

The dominant narrative about the benefits of robo-advising suggests that low fees allow advising many users that would otherwise be unadvised. And, indeed, the fees robo-advisors charge are substantially lower than those human advisors would typically charge. Consistently, Reher and Sokolinski (2019) document that robo-advising increases the share of households exposed to financial advice, especially in the middle-class segment. Because robo-advisors can manage many small accounts at low cost, common limitations to the take up of financial advice, such as minimum account balances and high fees, can be easily overcome. Robo-advising can thus help reduce wealth inequalities by allowing middle-class consumers to enjoy the higher returns and tax advantages that were typically reserved to high-wealth investors through (costly) human financial advice.

At the same time, the incentives to provide robo-advice by FinTech institutions might also contribute to increase inequalities. On the one hand, the quality of robo-advising services varies substantially with the wealth of investors. Wealthier investors, who are willing to pay higher fees and hence from which robo-advising companies can obtain higher margins, often enjoy more precisely tailored and better directed advice (e.g., see D'Acunto et al. 2019). On the other hand, low-income households, who often barely access financial institutions, who finance their spending with high-interest borrowing such as payday loans, and who make financial mistakes due to the lack of financial literacy, are perhaps the category that would need financial advice the most. And, yet, existing robo-advisors do not cater to this segment because of their limited wealth accumulation for retirement and hence the limited scope for fee extraction.

The lack of products catering to low-income households has two important implications. First, if anything, robo-advising might *increase* wealth inequality in the broader population, as middle-income households would increase their wealth and wealthy households would increase it by even more, but low-income households would not improve. Second, the question of who should provide robo-advising services to low-income households becomes prominent: Is there a business model that might allow private providers to target such population? Otherwise, should the public sector provide robo-advising for low-income households? After all, robo-advising for low-income households might replace costly debt-repayment programs. Answering these questions requires observational evidence or well-crafted randomized control

trials (RCTs) to evaluate public programs that provide cheap or free financial advice to low-income households.

4.4 Systemic Implications of Homogenizing Investors Through Robo-Advising

Robo-advisors might also have unpalatable systemic implications. For the case of human financial advice, tailoring advice to clients' needs and wants and the different preferences and beliefs of investors guarantee substantial differentiation in the cross section of portfolios. Most existing robo-advisors, instead, make a large number of users invest in the same exact portfolios. The largest portion of such portfolios is based on indexed ETFs. An economy in which a large part of households follows robo-advice could thus be much more exposed to the effects of aggregate negative shocks (e.g., see Bond and Garcia 2019). Whether indexing and robo-advising produce this higher sensitivity of aggregate wealth to business cycles and other shocks is an open question that is waiting for a setting that allows plausible counterfactuals.

4.5 What Ethical and Legal Standards for Robo-Advisors?

The last open question we discuss is the definition of ethical and legal standards for robo-advisors. We focus on three points. First, is the question of whether robo-advisors are fiduciaries. US regulators require that robo-advisors register under the Investment Advisers Act of 1940. Registration implies robo-advisors are fiduciaries. At the same time, though, robo-advisors display design features that cast doubts on whether they can truly act as fiduciaries, especially in terms of satisfying the duty of care (Fein 2017). Most robo-advisors do not provide holistic financial advice on all assets investors hold, but only on their financial portfolios. In fact, most robo-advisors do not even consider investors' asset holdings at other institutions when framing investment strategies. Moreover, robo-advisors use self-assessment questionnaires in terms of preferences for risk and other characteristics, which casts doubt on whether robo-advisors can perform the appropriate personalized due diligence that would fulfill their duty of care (Fein 2017).

On top of further developing the legal theory behind robo-advising, several questions are open for financial economists. For instance, how to detect potential biases and discrimination in robo-advising algorithms? Are new professional types needed, such as regulator/computer scientists, who

can bring together a strong legal background with an understanding of the mechanics of complex algorithms? Ultimately, how can regulators implement their assessments of whether fiduciary duties have been breached in a context in which the language in which the advice is delivered—the algorithms—are barely understandable to many legal practitioners?

A second issue is the unprecedented concentration of sensitive personal data that robo-advisors and their developers obtain. Holistic advice requires access to almost any feature of individuals' private sphere. Do robo-advising developers have the right to exploit such unique data commercially, either by selling the data to third parties or by elaborating them to provide consulting services? To what extent are investors aware of the importance of this issue, and can they form a meaningful assessment of the dollar value of these data given the large algorithmic illiteracy in the broader population? Tang (2019) makes progress on this question by exploiting a unique peer-to-peer lending setting to quantify the value users attach to privacy.

Even if robo-advisors excluded the possibility of selling or using data for purposes other than advice, increasingly frequent cyberattacks and data breaches would make the concentration of individual personal data in the hands of a few robo-advisors risky. Akey et al. (2020) estimate the value of unexpected data breaches in terms of corporate reputation and subsequent firm policies. One might also worry that the concentration of so much personal information about the broader population in the hands of a few providers and the vulnerability of such providers to internal and foreign cyberattacks might represent a matter of national security. Could this argument be developed further to support the recent proposals to break up big tech companies? Financial economists need to contribute to this debate by providing data and facts about the relevance of data breaches and the far-reaching consequences of data leakages, if any.

A third issue refers to the institutional contexts in which robo-advisors develop. If private corporations operate in competitive markets, information on the universe of individuals will not be concentrated, which reduces the potential damage of data leakages and the use of data for purposes other than providing robo-advising services. At the same time, private companies might be reluctant to share their data for public security purposes to avoid breaching the confidentiality of the information on their clients. Recent developments in operations research, such as Cai and Kou (2019), propose algorithms that allow statistical inference with encrypted data, thus guaranteeing individuals' anonymity.

Very different issues arise in settings in which governments control at the same time all major sources of information and means of production in the

economy. More theoretical and empirical research is needed to understand the political-economy implications of data concentration and lack of anonymity in these contexts.

References

- Akey, Pat, Vincent Gregoire, and Charles Martineau. 2020. "Price Revelation from Insider Trading: Evidence from Hacked Earning News" Working Paper available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3365024.
- Alsabah, Humoud, Agostino Capponi, Octavio R. Lacedelli, and Matt Stern. 2019. Robo-advising: Learning investors' risk preferences via portfolio choices. *Journal of Financial Econometrics*: forthcoming.
- Baker, Scott. 2018. Debt and the response to household income shocks: Validation and application to linked financial account data. *Journal of Political Economy* 126 (4): 1504–1557.
- Baker, Tom, and Benedict G. Dellaert. 2019. Behavioral finance, decumulation, and the regulatory strategy of robo-advice. In *The disruptive impact of FinTech on retirement systems*, ed. Julie Agnew and Olivia Mitchell. Oxford: Oxford University Press.
- Barber, Brad M., and Terrance Odean. 2001. Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics* 116: 261–92.
- Ben-David, Daniel, and Orly Sade. 2018. Robo-advisor adoption, willingness to pay, and trust—An empirical investigation. Mimeo.
- Ben-David, Daniel, Orly Sade, and Ido Mintz. 2019. Using AI and behavioral finance to cope with limited attention and reduce overdraft fees. Mimeo.
- Bhattacharya, Utpal, Andreas Hackethal, Simon Kaesler, Benjamin Loos, and Steffen Meyer. 2014. Is unbiased financial advice to retail investors sufficient? Answers from a Large Field Study. *Review of Financial Studies* 25 (4): 975–1032.
- Bond, Philip, and Diego Garcia. 2019. The equilibrium consequences of indexing. Mimeo.
- Browning, Martin, and Thomas F. Crossley. 2001. The life-cycle model of consumption and saving. *Journal of Economic Perspectives* 15 (3): 3–22.
- Cai, Ning, and Steven Kou. 2019. Econometrics with privacy preservation. *Operations Research* 67 (4): 905–926.
- Campbell, John Y. 2006. Household Finance. *The Journal of Finance* 61: 1553–1604.
- Carroll, Christopher. 1997. Buffer stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112 (1): 1–56.
- Carroll, Christopher. 2000. Requiem for the representative consumer? Aggregate implications of microeconomic consumption behavior. *American Economic Review* 90 (2): 110–115.

- Castelo, N., Maarten W. Bos, and Donald R. Lehmann. 2019. Task-dependent algorithm aversion. *Journal of Marketing Research* 56: 809–825.
- Cocco, Joao. 2005. Portfolio choice in the presence of housing. *Review of Financial Studies* 18 (2): 535–567.
- Cocco, Joao, and Francisco Gomes. 2005. Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18 (2): 491–533.
- D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi. 2019. The Promises and Pitfalls of Robo-advising. *Review of Financial Studies* 32(5): 1983–2020.
- D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2019a. IQ, expectations, and choice. NBER Working Paper No. 25496.
- D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2019b. Cognitive abilities and inflation expectations. *AEA Papers and Proceedings* 109: 562–566.
- D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi. 2019c. The promises and pitfalls of robo-advising. *Review of Financial Studies* 32 (5): 1938–2020.
- D'Acunto, Francesco, Marcel Prokopczuk, and Michael Weber. 2019d. Historical antisemitism, ethnic specialization, and financial development. *Review of Economic Studies* 86 (3): 1170–1206.
- D'Acunto, Francesco, Thomas Rauter, Christopher Scheuch, and Michael Weber. 2019e. Perceived precautionary savings motives: Evidence from FinTech. Mimeo.
- D'Acunto, Francesco, Alberto G. Rossi, and Michael Weber. 2019f. Crowdsourcing financial information to change spending behavior. Mimeo.
- Fein, Melanie. 2017. Are robo-advisors fiduciaries? Mimeo.
- Foerster, Stephen, Juhani T. Linnainmaa, Brian T. Melzer, and Alessandro Previtero. 2017. Retail Financial Advice: Does One Size Fit All? *The Journal of Finance* 72 (4): 1441–1482.
- French, Kenneth R., and James M. Poterba. 1991. Investor Diversification and International Equity Markets. *The American Economic Review* 81(2): 222–226.
- Gargano, Antonio, and Alberto Rossi. 2018. Does it pay to pay attention? *Review of Financial Studies* 31 (12): 4595–4649.
- Gargano, Antonio, and Alberto Rossi. 2019. There is an App for that: Goal setting and saving in the FinTech Era. Mimeo.
- Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny. 2015. Money doctors. *Journal of Finance* 70 (1): 91–114.
- Gomes, Francisco, Kenton Hoyem, Wei-Yin Hu, and Enrichetta Ravina. 2020. “Retirement Savings Adequacy in Retirement Contribution Plans.” Working paper available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3294422.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2008. Trusting the stock market. *Journal of Finance* 63 (6): 2557–2600.
- Hackethal, Andreas, Michael Haliassos, and Tullio Jappelli. 2012. Financial advisors: A case of babysitters? *Journal of Banking & Finance*, 36 (2): 509–524.

- Han, Bing, David Hirshleifer, and Johan Walden. 2019. Visibility bias in the transmission of consumption beliefs and undersaving. NBER Working Paper No. 25566.
- Hillis, A. 2017. Salience through information technology: The effect of balance availability on the smoothing of SNAP benefits. Mimeo.
- Jung, Dominik, and Christof Weinhardt. 2018. Robo-Advisors and financial decision inertia: How choice architecture helps to reduce inertia in financial planning tools. Mimeo.
- Laibson, David. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112 (2): 443–478.
- Lee, Sung K. 2019. FinTech nudges: Overspending messages and personal finance management. Mimeo.
- Levi, Yaron. 2019. Online information design and consumer behavior. Mimeo.
- Linnainmaa, Juhani, Brian Melzer, and Alessandro Previtero. 2021. The misguided beliefs of financial advisors. *Journal of Finance* 76(2): 587–621.
- Logg, Jennifer M., Julia A. Minson, and Don A. Moore. 2019. Algorithm appreciation: People prefer algorithmic to human judgement. *Organizational Behavior and Human Decision Processes* 151: 90–103.
- Lusardi, Annamaria, and Olivia S. Mitchell. 2007. Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education. *Business Economics* 42 (1): 35–44.
- Mitchell, Olivia, S., and Stephen P. Utkus. 2004. “Lessons From Behavioral Finance for Retirement Plan Design.” In O. S. Mitchell and S. P. Utkus, eds., *Pension Design and Structure: New Lessons from Behavioral Finance*. New York: Oxford University Press.
- Reher, Michael, and Stanislaw Sokolinski. 2019. Does FinTech democratize investing? Mimeo.
- Reher, Michael, and Celine Sun. 2019. Automated financial management: Diversification and account-size flexibility. *Journal of Investment Management* 17 (2): 1–13.
- Rossi, Alberto G., and Stephen Utkus. 2019a. Who benefits from robo-advising? Evidence from machine learning. Mimeo.
- Rossi, Alberto G., and Stephen Utkus. 2019b. The needs and wants in financial advice: Human vs. robo-advising. Mimeo.
- Sironi, Paolo. 2016. Fintech innovation: From robo-advisors to goal based investing and gamification. Chichester: Wiley.
- Tang, Huang. 2019. The value of privacy: Evidence from online borrowers. Mimeo.
- Terrance, Odean. 1998. Are Investors Reluctant to Realize Their Losses? *The Journal of Finance* 53 (5): 1775–1798.
- Terrance, Odean. 1999. Do Investors Trade Too Much? *American Economic Review* 89 (5): 1279–1298.



Implementation of Open Banking Protocols Around the World

Tania Ziegler

1 Introduction

Given its somewhat amorphous meaning as both a regulatory initiative and financial technology solution, I first define Open Banking and the key stakeholders engaged with it: data attribute providers, third-party providers (TPPs), and customers. When considering how these different stakeholders interact with one another, the concept of Open Banking as an Open API helps inform the different approaches taken by regulators when implementing a standard protocol. At present, there exists no globally harmonised set of rules related to API protocol. In the most general sense, API standards address three overarching themes: that of data standards, specification standards, and security standards. In this chapter, I review the various standards in place, especially as they influence a given jurisdiction's Open Banking implementation goals.

As a particular example, I then review the development and implementation of the United Kingdom (UK) Open Banking standard. The UK example serves as a pioneering example of Open Banking, and one that many other jurisdictions refer to when forming their own Open Banking approach. Hence I use the UK example as a road map to discuss and review

T. Ziegler (✉)

Cambridge Judge Business School, University of Cambridge, Cambridge, UK
e-mail: t.ziegler@jbs.cam.ac.uk

several additional regulatory case-studies. The UK example also serves as a useful explanation of how Open Banking regulation can respond to legislative mandates and the evolution of thinking towards data ownership, sharing, and use.

As regulators introduce their own Open Banking interventions, it is important to understand the different statutory objectives which dictate a given regulator's approach. The primary objective underlying the regulatory authority will heavily influence that jurisdiction's specific Open Banking guidelines in terms of who the stakeholders are, the relationship they have with data, and if formal supervision will take place. The three statutory objectives under review are Competition, Market Integrity, and Consumer Protection.

I then examine five additional regulatory case-studies (Australia, Hong Kong, Singapore, Brazil and Mexico) to provide a comprehensive overview of how Open Banking is developing on the world stage. To inform this discussion, I review the underlying statutory objectives that influence each jurisdiction's Open Banking intervention to distinguish between the different flavours of Open Banking and to evaluate the key tendencies of each approach. I do not seek to qualify which approach is superior to another. Throughout each country case-study, the UK example will be used to serve as a reference and point of comparison.

1.1 Defining Open Banking

With increased digitalisation of financial instruments and channels in recent years, it comes as no surprise that the application of a financial technology solution is now aimed at the traditional retail banking industry.

In recent years, Open Banking has captured the attention of the financial world, with regulators and policymakers making determinations on how best to implement Open Banking rules; and Banks, FinTechs and other financial service providers trying to determine how best to utilize Open Banking to maximise their business model while not losing their stake in the marketplace.

A somewhat amorphous term, Open Banking can be defined as a collaborative financial technology programme that is aimed at creating a standard for data sharing protocols via Application Programming Interfaces (APIs)¹ between two or more unaffiliated parties, in order to deliver enhanced capabilities to the incumbent marketplace.

¹ <https://www.openbanking.org.uk/wp-content/uploads/open-banking-report-150719.pdf>.

Directed at rebalancing retail banking, the concept of Open Banking seeks to enable increased competition and customer choice, thus disrupting the incumbent banking sector and empowering individual and business consumers to own and dictate the use of their financial data. By utilising technology to rebalance markets in favour of these consumer users, such initiatives allow bank customers to securely share their data to third-party providers (TPPs), in particular non-bank financial services providers. By doing so, customers can take advantage of greater competition across the banking sector, taking advantage of comprehensive and better access to financial services, lower prices, and greater choice. In effect, Open Banking addresses three critical areas; money management, payment processing, and lending.

At the crux of the Open Banking debate is *ownership of data*. At a fundamental level, Open Banking changes the relationship between banks and the client data they hold. Open Banking quite specifically aims at shifting the ownership of consumer data, where banks move from data-owners to data-custodians. Traditionally, a concentrated number of banks have held and controlled underlying client data, imposing considerable limitations on how that data can be used or shared outside of their own organisation. An Open Banking standard shifts this dynamic by placing this data back into the consumer or business client's ownership. Accordingly, the Open Data Institute has outlined the three types of stakeholders or participants engaged by Open Banking.

- Data attribute providers: banks, financial services companies, and other organisations through whom data is stored and shared;
- Third parties: developers, FinTech, and other organisations who use data provided to design and offer new products; and
- Customers: individuals and businesses who share their data; publishers of open data.²

Open Banking changes the way that these participants engage with one another and their relationship with client data. In the first instance, the 'data attribute provider' is the custodial organisation, holding or storing data on behalf of their customers. In layman's terms, this would be a bank, who previously was treated as the 'owner' of the data, not its steward. The third party

² The Open Banking standard: Unlocking the potential of Open Banking to improve competition, efficiency and stimulate innovation, prepared by the Open Banking Working Group, 16 February 2016. <https://www.paymentsforum.uk/sites/default/files/documents/Background%20Document%20No.%202%20-%20The%20Open%20Banking%20Standard%20-%20Full%20Report.pdf> (p. 12).

(TPP), in the form of an external organisation, hopes to access the data held by the ‘data attribute provider’. The customer, an individual or business, is now viewed as the owner of the data held by the ‘data attribute provider’, and it is at their discretion how that data may be used by a TPP. In this case, the customer determines if, how, and for how long their data can be used by a TPP. The customer now has the authority to permit, refuse, or revoke access to their data, using an opt-in approach. This data sharing relationship is fundamentally different and may enable additional value-added products and services for the customer.

With this in mind, Open Banking places obligations on how such data can be held, and implements protocols as it relates to the sharing of data. Data sharing, both in how it is stored/collected and used, has three overarching technical principles as related to an individual’s rights to their data:

- data portability—the individual may share their data freely with whomever they choose;
- consent—the individual must provide explicit consent to sharing their data;
- specific usage—the individual’s data may only be used for the pre-agreed purposes.³

As a technological solution, the concept of an Open Banking standard has become somewhat synonymous to Open APIs. It is generally accepted that such an API approach needs to align with certain standardisations and specifications, though globally harmonised rules have yet to exist. Nevertheless, the UK’s Open Banking Standard address three overarching themes when considering an open API environment for Open Banking. These standards set the parameters in which an API specification can deliver account information. They are:

- Data Standards—Rules by which data is described and recorded (i.e. is the data formatted and defined the same way across different data attribute providers);
- API Standards—Specifications to inform the design, development, and maintenance of an open API (i.e. the technical parameters imposed on an API);

³ The Open Banking standard: Unlocking the potential of Open Banking to improve competition, efficiency and stimulate innovation, prepared by the Open Banking Working Group, 16 February 2016. <https://www.paymentsforum.uk/sites/default/files/documents/Background%20Document%20No.%202%20-%20The%20Open%20Banking%20Standard%20-%20Full%20Report.pdf> (p. 14).

- Security Standards—Security aspects of the API specification (i.e. to hedge cyber security risks, risks associated to individual or personal data, etc.).⁴

Though not universally mandated by a regulator, some iteration of these three protocols appear in most regulatory approaches to Open Banking and I will refer to them when reviewing country-level case-studies.

2 An Overview of Open Banking in the UK

In recent years, Open Banking often also refers to a series of regulatory and policy initiatives aimed at rebalancing retail banking. Although considerable technical considerations are at play, especially as related to the implementation of Open APIs, Open Banking as a regulatory concept is critical. This is because it is a regulatory body which will ultimately be tasked with the standard setting necessary to supervise and oversee an Open Banking schema.

Potential opportunities of Open Banking initiatives include improved pricing transparency in personal and business current accounts; improved transparency about the quality of current accounts; reduced switching costs between service providers; and lower barriers to entry for new service providers. Equally, regulatory authorities must balance these potential opportunities against several challenges that might arise in implementation, including: setting key technical and user standards; determining minimally viable operational guidelines; deciding upon which authorised entities can participate; ensuring consumer protection and dispute resolution; and effective monitoring of new API protocols.

As such, this chapter will frame the regulatory approach to Open Banking initiative as developed in the United Kingdom. Though there are ample examples of other jurisdictions that have implemented an Open Banking initiative, none are as far advanced nor well documented as that of the United Kingdom.

The UK's iteration of Open Banking is a response to a proactive EU initiative to legislate rules for data sharing. As described by Cortet et al., the updated Payment Services Directive (PSD2) is not simply

⁴ Adopted from page 9, The Open Banking standard: Unlocking the potential of Open Banking to improve competition, efficiency and stimulate innovation, prepared by the Open Banking Working Group, 16 February 2016. <https://www.paymentsforum.uk/sites/default/files/documents/Bacground%20Document%20No.%202%20-%20The%20Open%20Banking%20Standard%20-%20Full%20Report.pdf>.

another regulation requiring a mere operational and compliance approach, but an accelerator of the already on-going change in the fabric of the — digitisation prone — financial industry. The PSD2 provisions on ‘Access to account’ for Payment Initiation and Account Information Services (‘XS2A’) will accelerate this by forcing banks to open up consumer payment accounts for appropriately licensed, innovative (bank and non-bank FinTech) service providers.⁵

2.1 PSD2 as a Framework for Open Banking

Since 2013, the European Commission had identified revisions to existing data sharing rules as a critical component of reform to its Payments Services Directive. The origins of Open Banking in the UK can be traced back to the 2014 Fingleton Report, commissioned by HM Treasury and undertaken by the Open Data Institute (ODI) and Fingleton Associates. The purpose of this paper was to assess the opportunities for improving UK banking, especially in light of the failed 2011 Midata initiative. Though this initiative marked the first attempts with the UK to move towards Open Banking, flaws in the design and implementation rendered this attempt a failure. The Fingleton Report concluded that a policy initiative (such as that of Midata), would ‘benefit from employing more common technology and standards for data sharing. Common standards were identified as important for interoperability between different providers and to prevent incumbents from developing ‘walled gardens’ that kept their customers from accessing rivals’ products.’⁶

In 2015, HM Treasury established the Open Banking Working Group (OBWG), tasking this body to determine how data sharing might work practically. The following year, and alongside the release of PSD2, this working group published guidance and high-level recommendations on how data sharing could feasibly look, using standardised APIs. Critically, the OBWG recommended that ‘Open Banking adopt a decentralised system across different banks, which would be more secure than a single, centralised system.’⁷

Simultaneously, by December of 2015,⁸ a final version of PSD2 was published, setting ‘requirements for Payment Account Providers to allow third parties – with appropriate consent – to share account information and

⁵ Cortet, Mounaim, Tom Rijks, Shikko Nijland. 2016, Spring. PSD2: The digital transformation accelerator for banks. *Journal of Payments Strategy & Systems* 10 (1): 13–27. Henry Stewart Publications.

⁶ <https://www.openbanking.org.uk/wp-content/uploads/open-banking-report-150719.pdf> (p. 13).

⁷ *Ibid.*, (p. 14).

⁸ https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:JOL_2015_337_R_0002&from=EN.

to initiate payments.⁹ Essentially, this directive set the rules of engagement by mandating that banks present their data in the form of an API. Unlike most directives, PSD2 required member states to comply with PSD2 within two years of publication (in this case, January 2018), and set its basic provisions into national law. As such, PSD2 can widely be recognised as the framework which Open Banking is based upon. In a 2018 Wired article, Manthorpe makes the following distinction between Open Banking and PSD2:

Sometimes the two get confused: essentially, Open Banking is the UK version of PSD2. The difference is that whereas PSD2 requires banks to open up their data to third parties, Open Banking dictates that they do so in a standard format.¹⁰

While PSD2 sets out the basic rules of engagement vis a vis banking data, the Open Banking initiative sought to deliver these changes. It is not uncommon for Open Banking and PSD2 to be conflated. Yet it is important to ensure that certain distinctions are understood. The first is that PSD2 covers all payment accounts and methods (credit cards, prepaid cards, e-wallets), while Open Banking in the UK is limited to personal and business accounts. Yet, in contrast to PSD2, the Open Banking initiative in the UK has been more explicit around the definition and development of the required APIs, as well as the security and messaging standards.¹¹

Following the Retail Banking Market Investigation Order 2017 by the Competition and Markets authority (CMA), the CMA initiated the first stage of Open Banking within the UK to ‘stimulate innovation across the financial sector, to enable the ‘unbundling’ of complex retail banking products (in particular personal and business current accounts), and to ultimately lead to greater competition overall.’¹² As a result of this order, the Open Banking Implementation Entity¹³ (OBIE) was created to create software standards and industry guidelines that drive competition and innovation in UK retail

⁹ <https://www.paymentsforum.uk/sites/default/files/documents/Background%20Document%20No.%202020-%20The%20Open%20Banking%20Standard%20-%20Full%20Report.pdf> (p. 14).

¹⁰ What is Open Banking and PSD2? WIRED explains; Manthorpe, Rowland, 17 April 2018. <https://www.wired.co.uk/article/open-banking-cma-psd2-explained>.

¹¹ Markos Zachariadis, Pinar Ozcan, The API Open Banking, 15 June 2017, SWIFT Institute Working Paper No 2016 - 001 (p. 4).

¹² CMA, Retail banking market investigation: Final report. 2016. <https://www.gov.uk/cma-cases/review-of-banking-for-small-and-medium-sized-businesses-smes-in-the-uk>.

¹³ Funding for the OBIE comes from the CMA9, while the CMA, the Financial Conduct Authority (FCA), and HMT provide governance oversight.

banking in September 2016.¹⁴ The OBIE serves as the central implementation entity tasked with writing standards, building supporting infrastructure and overseeing the Open Banking implementation across the mandated and obligated banks. Critically, ‘the OBIE has been structured to combine oversight and monitoring of the banks with implementation support provided to both banks and TPPs.’¹⁵

In line with the PSD2 rollout, the first phase of the Open Banking implementation began in 2018, and was completed in September of 2019. In the first instance, the CMA ordered that the largest banks within Great Britain and Northern Ireland (known as the CMA 9), were required to allow for their personal and business customers to access and share their data to authorised third parties by the 2018 deadline.¹⁶ The order also indicated that the use of secured APIs be used to make this data sharing possible, in a standardised format. The first iteration of Open Banking APIs is often referred to as the Open Banking APIs Version 1.

With its second year now ending, Open Banking in the UK has achieved some impressive milestones. By the end of 2019, there are now 204 regulated providers (70 account providers, 134 TPPs, 61 regulated entities) and has recorded 1.25 billion API calls within its ecosystem.¹⁷ It should be noted that every provider must be regulated by the Financial Conduct Authority (FCA) or their European equivalent.

Within 2020, it is the goal of the OBIE to finish their implementation of the Open Banking framework, and to further improve usability and functionality for consumer stakeholders and TPPs. For instance, at the end of December 2019, the OBIE put forth further updates to its API specifications, customer experience guidelines, and operational guidelines (referred together as Open Banking Standard Version 3.1.4). In effect, the OBIE is treating these guidelines as something of a living document, with innovation and stakeholder usability at its core.

Looking at the future of Open Banking in the UK, it is worth noting that the FCA has created an Advisory Group in order to drive forward an ‘open finance’ strategy. This strategy would extend beyond Open Banking, and refers to data sharing of a wider range of financial products beyond traditional banking data. As outlined in the FCA’s 2019/2020 Business Plan,¹⁸

¹⁴ <https://www.openbanking.org.uk/about-us/>.

¹⁵ <https://www.openbanking.org.uk/wp-content/uploads/open-banking-report-150719.pdf> (p. 35).

¹⁶ These banks include: Lloyds, Barclays, Nationwide, RBS, Santander, Danske Bank, HSBC, Allied Irish Banks and Bank of Ireland.

¹⁷ <https://www.openbanking.org.uk/wp-content/uploads/2019-Highlights.pdf>.

¹⁸ <https://www.fca.org.uk/publication/business-plans/business-plan-2019-20.pdf>.

the advisory group ‘will review the effectiveness of Open Banking while also leading the broader public debate on Open Finance, including seeking out opportunities to use this infrastructure to foster competition in the savings sector.’¹⁹

3 Regulatory Objectives and How They Dictate Open Banking

The Open Banking phenomena has captured the attention of policymakers and regulators globally, with many countries and inter-governmental organisations exploring how such programmes might be pursued within their own jurisdictions. When considering a country’s approach to Open Banking, it is not surprising that such a schema would align with the underlying statutory objectives that guide the overseeing regulatory authority tasked with implementation.

By reviewing which key statutory objectives are being prioritised by the corresponding regulating authority, we can review how Open Banking initiatives are evolving globally, and consider the local context dictating a given jurisdiction’s unique approach.

If our assumption is that the Open Banking phenomena and its particularly brand within a country will align with the regulators existing and/or emerging priorities, then it is worth outlining the potential statutory objectives prioritised by regulatory authorities in key case-study countries. This chapter will continue to use the UK example of Open Banking as an explanatory road map, while reviewing additional Open Banking interventions from an international perspective; Australia, Singapore, Hong Kong, Mexico, and Brazil.

When considering the factors that a regulator must take into account when performing its main functions, the statutory objectives outlined by its governing legislation are of paramount importance. Though other factors may play an important role when considering its regulatory principles, stated operational objectives will dictate how a regulator approaches its supervisory mandate vis a vis relevant market functions.

When considering financial services and markets, three operational statutory objectives seem to hold constant across the case-study countries, but to varying degrees of importance. They are *Competition*, *Consumer Protection*, and *Market Integrity*.

¹⁹ Ibid., (p. 22).

3.1 Placing Competition at the Heart of Open Banking Implementation

In the broadest sense, competition can be defined as ‘a process of rivalry between suppliers that takes place either in or for the market.’²⁰ When related to financial markets, this objective is rooted in promoting effective competition in the interest of consumers within the financial markets. It relates to the needs of different consumers who may use services, including their need for information that enables informed decision making. It also relates to the ease which consumer may obtain, change, or withdraw from using services. Critically, competition can effectively promote innovation of new products and services.

On its face, competition as an objective, can seem somewhat counter-intuitive to regulatory interventions. In its January 2020 CMA 111 review of Regulation and Competition, the CMA acknowledged that ‘there is a concern that regulation can have the effect of stifling competition, and thereby deprive customers [...] through raising barriers to entry.’²¹ The report goes on to say that:

In order to get the appropriate balance between the two, i.e. to ensure that the level of regulation is proportionate and does not impose any unnecessary restrictions on competition beyond securing specific policy objectives, it is important that there is an on-going focus on the process of designing and implementing ‘better regulation’ across government and the public sector more generally.²²

The Open Banking policy objective, at its core, is a FinTech intervention. In this context, it is not surprising that a key existential question that regulators often combat when thinking about competition and regulation is whether to regulate before a technology or market has entered its economy or take a ‘wait & see’ approach. To contextualise this, a discussion on ex-ante versus ex-post regulation is useful.

Ex-ante regulation indicates that regulators will implement preventive measures to lower the risk related to FinTech. However, such measures may deter the innovation and growth of industry, and ipso facto competition.

²⁰ Regulation & competition: A review of the evidence, a review of the evidence on the impact of regulation on competition in the UK; Competition & Markets Authority, 10 January 2020, CMA 111. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/857024/Regulation_and_Competition_report_-_web_version.pdf.

²¹ Ibid., (p. 14).

²² Ibid.

On the other hand, Ex-post regulation indicates that regulators implement measures when risks related to FinTech materialise. Such measures may promote the growth of the market, and the resultant proliferation of products and services, but the risk may be greater than ex-ante regulation.

With this consideration, according to CMA 111, an ongoing national and international debate around this very issue persists, and is 'focused in part on how to promote competition and on the role of ex-ante regulation [... concluding that markets with existing dominant players are] unlikely to self-correct. As a result, at the very least competition policy needs to be adapted/updated to properly. [In the UK context, the regulator] would have a remit to develop ex-ante regulatory tools and frameworks to support greater competition and consumer choice.'²³

When we consider the Open Banking initiative in the UK, this intervention serves as 'an example of a regulatory remedy that was designed to ensure that customers benefitted from technological advances and that new entrants and smaller providers in the UK retail banking sector were able to compete more fairly.'²⁴ What has made the UK's Open Banking model quite unique and rooted in a competition mandate is its creation of the OPIE. As noted earlier, the OPIE is mandated by the CMA but comprised of stakeholders from across the traditional banking sector as well as new and emerging FinTechs. Multiple regulatory bodies also feed into this entity (such as the FCA), influencing the overarching regulatory approach.

We have established, then, that the CMA's approach via the OBIE is that of an ex-post approach, rooted in three essential prescriptions. They are:

Principles-based regulation: this approach entails moving away from a reliance on detailed, prescriptive rules and relies instead on high-level, broadly stated principles to set the standards by which regulated firms must conduct business. This leaves firms with the flexibility to determine how they comply with those principles.

Codes of conduct: a form of regulation that applies only to firms in an industry that satisfy certain criteria. The code of conduct will set certain restrictions on the behaviour of these firms, for example how they must treat their suppliers; but it will give the regulated firms some discretion in how they comply with the code. A code can also be changed with industry agreement as circumstances change.

Participative regulation: regulation in which there is a greater degree of engagement between firms and the regulator in a market, with firms making formal proposals to the regulator e.g. in relation to the introduction of new

²³ Ibid., (pp. 15–16).

²⁴ Ibid., (p. 15).

services or products. This can be particularly helpful for new entrants wanting to bring products to a market, particularly if the regulator can then forbear from regulating, until there is a better sense of whether intervention is needed / what form that intervention should take.²⁵

When reviewed on balance, a key facet of the OBEI's approach to implementation has been the MVP model, and continuous and flexible evaluations impacting rules and guidance as related to firms falling under an Open Banking standard (i.e. making use of the Open API structure).

Furthermore, it is only once the implementation phase is complete that the OBEI would transition towards a monitoring role, to ensure service levels are maintained and stakeholders continue meeting obligations. In the CMA 111, it is made clear that 'the benefits of ex-post evaluations and case study evidence clearly points to the importance of evaluating market interventions through Post Implementation Reviews, and of these focusing on competition impacts.'²⁶

3.2 Competition Policy as an Innovation Driver

Innovation, though in and of itself not a statutory objective, can be viewed as an emerging non-statutory priority for regulators and policymakers introducing an Open Banking initiative. A call for innovation from policymakers may even serve as a proxy for an explicit competition mandate in jurisdictions that lack this as a legislated objective. As Milne explains, banking regulation is largely rooted in prudential requirements and consumer protection, and not the prescriptive promotion of competition and innovation. When considering incumbent market players, there exist 'weak incentives for innovation, [thus explaining why] most long-established banks continue to struggle with managing and maintaining a huge body of legacy software, rather than replacing with integrated systems.'²⁷

Cue Open Banking and its competition policy! As Milne notes, competition policy can lead to innovation in network industries, such as retail banking. He explains that 'competition interventions have until recently been based on the application of general competition law, rather than on specific regulatory frameworks tailored to the specific network features of banking

²⁵ Ibid., (pp. 70–71).

²⁶ Ibid., (p. 11).

²⁷ Milne, A. 2016. Competition policy and the financial technology revolution in banking. *DigiWorld Economic Journal* 103: 145–161.

and other financial services.²⁸ Explicit market interventions, such as Open Banking imposed on incumbent banks, flips this conception on its head and suggests that competition-focused regulation may address banking-sector innovation. He notes, ‘authorities are now beginning to address seriously the challenge of regulating access to banking platforms, addressing more fully the implications of banking networks for competition and innovation, a key step if the full potential benefits of FinTech are to be realised.’²⁹

As suggested by Cortet et al., regulatory interventions can accelerate technological innovation by ‘forcing banks to use APIs to open up consumer payment accounts [...] potentially presenting a strategic challenge to incumbent banks, who will have to focus on ensuring they exploit the opportunities of technology to improve customer experience.’³⁰ With respect to network effects, the proactive nature of incumbents will dictate both opportunity, or risk. By imposing a mandate for all incumbent banks to adhere to an ‘Open Banking API’, incumbents can either view this as an opportunity to position themselves as innovators alongside TPPs, serving their clients in an enhanced capacity. Alternatively, by limiting themselves to bare-minimum thresholds of compliance, they may risk loss of market share and customers to more agile TPPs.

3.2.1 The Role of a Consumer Protection Objective

When we think of the retail banking sector, consumer protection comes to mind. As a statutory objective, this means that consumers (individuals or business customers) receive appropriate levels of protections alongside their activities. A regulator may determine differing degrees of protections or safeguards related to a consumer profile. In the case of financial markets, protections include general principles regarding decision making, appropriateness-standards, and access to ‘fair, true and not misleading’ information.

Alongside the OBIE’s competition considerations, concurrent considerations regarding consumer protection exist when implementing Open Banking. This means that the range of potential regulatory interventions has increased, with the CMA and additional sectoral regulators (in this case, the FCA³¹) incorporating more exacting regulatory interventions to

²⁸ Ibid.

²⁹ Ibid.

³⁰ Cortet, Mounaim, Tom Rijks, Shikko Nijland. 2016, Spring. PSD2: The digital transformation accelerator for banks. *Journal of Payments Strategy & Systems* 15: 13–27.

³¹ It should be noted that the FCA’s primary statutory objective is Consumer Protection.

incorporate consumer protection objectives. It is worth highlighting that in July of 2019, the CMA and the FCA revised and recommitted to two critical MOUs, outlining shared powers as related to their primary statutory objective; concurrent competition powers³²; and concurrent consumer protection³³ powers. These two MOUs set out the parameters for shared powers and responsibilities as related to rule setting, supervision, and regulatory intervention. As such, a secondary yet still crucial statutory objective as related to the UK's Open Banking standard is that of consumer protection.

In order to ensure robust consumer protection, the authorising bod(ies) must ensure that sufficient safeguards are in place. This may require additional regulation. As noted in the previous section, this additional regulation may cause frictions to the superseding goal of competition. With any market innovation, it is possible that new forms of consumer detriment may present themselves. As such, regulatory intervention throughout the implementation phase may be necessary to prevent such situations from arising. In the case of the OBIE, their 'test and evaluate' approach has allowed for data providers (banks) and TPPs to test programmes and potential use-cases within a ring-fenced or isolated environment, under the supervision and monitoring of the appropriate regulator. In this case, the UK's approach attempts to mitigate consumer protection risk by mimicking something of a 'product-testing' sandbox, creating a safe zone to allow innovators to live test products prior to formal regulation or compliance requirements. This sandbox approach gives the regulating body the opportunity to identify appropriate consumer protection safeguards as they are built into (1) new products and services (in the case of the TPP) or (2) API specifications, customer experience guidelines, and operational guidelines (in the case of the OBIE).

In the UK case-study, the role of PSD2 cannot be overlooked. This EU legislation is rooted in Consumer Protection considerations, with this directive enhancing consumer protection and associated consumer rights for member states. As such, the OBIE 'does not have complete discretion in its design of Open Banking. It has to fit within the EU's Second Payment Services Directive (PSD2).'³⁴ Where the CMA's dictating statutory objective

³² Memorandum of understanding between the Competition and Markets Authority and the Financial Conduct Authority—concurrent competition Powers (July 2019). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/823785/FCA_CMA_Competition_MOU.pdf.

³³ Memorandum of understanding between the competition and markets authority and the financial conduct authority on the use of concurrent powers under consumer protection legislation (July 2019). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/823786/FCA_CMA_-_MoU_consumer.pdf.

³⁴ <https://www.openbanking.org.uk/wp-content/uploads/open-banking-report-150719.pdf> (p. 4).

is to enable greater competition, the Open Banking implementation has to fit within the PSD2 framework and, along with it, a consumer protection statutory obligation.

As such, the UK's Open Banking schema places consumer protection front and centre, by embracing existing authorization and consent rules underpinned by PSD2. Specifically, this includes consumer security and retention of data control.

3.2.2 Market Integrity

Market integrity is the foundation of a fair and efficient financial market, ensuring that market participants and stakeholders enjoy equal access to markets, that price discovery and trading practices are fair, and that high standards of corporate governance exist. Perhaps the broadest of statutory objectives, *Market Integrity* relates to the active protecting and enhancing of the integrity of the financial systems; including its soundness, stability, and resilience.³⁵

As it relates to the UK's Open Banking approach, market integrity is not a key statutory objective. Yet, it is important to discuss this objective as it serves as the keystone for several international Open Banking approaches. As such, this statutory objective will not use the UK's example as a road map, but rather address the key elements as related to regulatory frameworks that emphasise market integrity.

Regulation with the aim of promoting market integrity places an emphasis on systemic stability and enhanced governance and conduct within the broader financial markets landscape. Regulatory interventions 'span financial and non-financial risk management (including cyber risk): conduct of business, investor and data protection, prudential requirements, and fitness requirements for individual risk-takers and market participants.'³⁶ With an emphasis on systemic stability, it is not all together surprising that regulatory regimes with a Market Integrity mandate focus on managing and minimising systemic risks associated with their supervisory sector. As it relates to financial markets, regulatory reforms may include 'enhanced oversight in line with [the] status of systemically important financial institutions [...] improv[ing]

³⁵ Market infrastructures and market integrity: A post-crisis journey and a vision for the future; Copyright 2018 Oliver Wyman and World Federation of Exchanges. https://www.world-exchanges.org/storage/app/media/research/Studies_Reports/wfe-amp-oliver-wyman-market-integrity-report.pdf.

³⁶ Ibid., (p. 16).



Fig. 1 International perspectives on Open Banking regulation and implementation

the transparency of risks by concentrating activity in central clearing counterparties, and to enhance resilience planning for periods of market stress' (Fig. 1).³⁷

The UK is the first country to develop an Open Banking initiative, with many other countries and territorial jurisdictions following its example. Though at various stages of development, the following country case-studies provide a good overview of how Open Banking is developing on the global arena. By understanding the underlying statutory objectives that influence the regulatory approach to Open Banking, we can begin to distinguish the different flavours of Open Banking arising on the world stage, and potentially evaluate their relative success and utility.

Australia

Open Banking in Australia looks most like that of the UK, though with some distinctive features. In August 2019, the Australian parliament passed Open Banking legislation. This legislation will make it mandatory for the largest four banks³⁸ in Australia to comply to consumer data sharing protocols (via

³⁷ Ibid., (p. 16).

³⁸ Since July 2019, CommBank, NAB, Westpac and ANZ were asked to make product data available via an open API, though this was not compulsory. July 2019 can be viewed as the starting phase of Open Banking in Australia, though binding legislation was not passed until the following month.

an Open API ecosystem) by July 2020.³⁹ Coinciding with new data sharing mandates is that of the newly legislated Consumer Data Right (CDR) system, which gives consumers regulated access to and control over the data held by service providers.⁴⁰ In its own words:

The CDR will give consumers greater access to and control over their data. It will improve consumers' ability to compare and switch between products and services, and will encourage competition between service providers, leading not only to better prices for customers but also more innovative products and services.⁴¹

Perhaps most comparable to the EU's PSD2 + GDPR,⁴² the CDR 'aims to provide greater choice and control for Australians over how their data is used and disclosed. Banking will be the first sector to which the CDR applies [...] to encourage greater competition.'⁴³ This will be followed by energy data and telecommunications data, aimed at creating greater data interoperability across different sectors. This goes beyond the scope of the UK's own Open Banking data sharing activities. Another key distinguishing factor is that the Australian Open Banking implementation is focused on data only, excluding payments/fund transfer activities.

Not surprisingly, the Australian Open Banking initiative takes many of its cues from the CDR. The CDR is regulated by both the Australian Competition and Consumer Commission (ACCC) and the Office of the Australian Information Commissioner (OAIC) 'as it concerns both competition and consumer matters as well as the privacy and confidentiality of consumer data. The ACCC leads on issues concerning the designation of new sectors of the economy to be subject to the CDR and the establishment of the CDR rules. The OAIC leads on matters relating to the protection of individual and small business consumer participants' privacy and confidentiality, and compliance with the CDR Privacy Safeguards'⁴⁴ The regulatory body tasked with implementing the Open Banking initiative is the ACCC, whose primary statutory objective of is to promote competition, though with an added Consumer

³⁹ <https://www.globalgovernmentforum.com/australia-delays-launch-of-open-banking-regime/>.

⁴⁰ Ibid.

⁴¹ <https://www.accc.gov.au/focus-areas/consumer-data-right-cdr-0>.

⁴² GDPR, or the General Data Protection Regulation is an EU directive that legislates protective standards on personal data, ensuring that a consumer has final say on how and where their data can be used. GDPR is often included in discussions of PSD2 as far as it relates to personal data ownership.

⁴³ Ibid.

⁴⁴ Ibid.

Protection element as relates to the CDR.⁴⁵ Their approach, like that of the UK, has included mandatory/compulsory orders to incumbent banks, formal regulation and registration of TPPs and related FinTechs (as accredited data recipients), governance frameworks, and API standardisation. Akin to the UK's MVP approach towards API standards, 'the ACCC has been looking for FinTech guinea pigs that will become accredited data recipients and help test the operations ecosystem surrounding the CDR legislation.'⁴⁶ It is worth noting, however, that Australia's API functionality is somewhat limited against that put forward by the UK. Specifically, read-only access is currently in play, limiting payment initiation or write-access.

Australian Open Banking will include several phase/roll-out periods, slowly adding additional types of banking and financial data from 'data attribute providers' beyond the big four banks. It is suggested that by 2021, every bank in Australia must provide product, account, and transaction data, including data on personal loans and mortgages.⁴⁷ Additionally, TPPs that participate in Open Banking will also be obliged to share their customer data, unlike UK Open Banking.⁴⁸ Unlike the UK, upon final phase out shareable data in Australia will include a broader scope of financial accounts, including lending and mortgage accounts.⁴⁹

Hong Kong

The Hong Kong Monetary Authority (HKMA) is the responsible regulatory body charged with implementing an Open Banking regime. There are considerable differences from the UK's own Open Banking, stemming primarily from overarching statutory objectives and priorities of the HKMA.

The HKMA is Hong Kong's central banking institution, with four main functions: maintaining currency stability within the framework of the Linked Exchange Rate System; promoting the stability and integrity of the financial system, including the banking system; helping to maintain Hong Kong's status as an international financial centre, including the maintenance and development of Hong Kong's financial infrastructure; and managing the Exchange Fund.⁵⁰ As it relates to Open Banking, this is most aligned with

⁴⁵ <https://www.accc.gov.au/about-us/australian-competition-consumer-commission/about-the-accc>.

⁴⁶ <https://www.computerweekly.com/news/252471528/Australia-gets-ready-for-open-banking>.

⁴⁷ <https://www.finder.com.au/open-banking>. .

⁴⁸ https://bankingblog.accenture.com/open-banking-framework-comes-to-australia?lang=en_US.

⁴⁹ Ibid.

⁵⁰ https://www.hkma.gov.hk/media/eng/publication-and-research/annual-report/2017/05c_About_the_HKMA.PDF.

objectives related to Market Integrity. Ultimately the aim is to develop entire ecosystems, a goal that is enabled by sharing data.

In HK, Open Banking is synonymous with its Open API framework,⁵¹ as overseen by the HKMA. Having launched phase 1 of a four phase API framework in January 2019, banks have been racing to publish APIs for a group of different financial products and services. However, unlike other Open Banking regimes which include compulsory rules for incumbent banks to adhere to data sharing, the HKMA has no such mandate. Rather, 'traditional retail banks are being encouraged to develop more personalized and novel services in collaboration with TPPs'⁵² in an effort to spur innovation. The four phases for the Open API are:

1. Product information: Deposit rates, credit card offerings, service charges, and other public information
2. Customer acquisition: New applications for credit cards, loans, and other products
3. Account information: Account balance, credit card outstanding balance, transaction records, credit limit change, and others
4. Transactions: Payments and Transfers⁵³

Coined as a 'supportive regulatory environment', the HKMA's role vis a vis Open Banking is to be more akin to an observer than active participant. For phases 1 and 2, there is no mandatory compulsion of retail banks, nor guidance or rules on API standardisation. Yet, now a year on, adoption of an Open API framework is on track and growing. According to the HKMA, in the first phase 20 retail banks made 500+ Open APIs available. Now entering its next phase, the Deputy Chief Executive of the HKMA, Howard Lee, has stated that '*as we move on to the next phase, with more sensitive data and complex functions covered, it is even more important to ensure security and safeguard customer*

⁵¹ <https://www.hkma.gov.hk/eng/key-functions/international-financial-centre/open-api-for-banking-sector/phased-approach.shtml>.

⁵² https://www.accenture.com/_acnmedia/pdf-99/centre-time-is-now-open-banking-hong-kong.pdf (p. 3).

⁵³ Adopted from open API phases as outlined on the HKMA website; The four phases of Open API. <https://www.hkma.gov.hk/eng/key-functions/international-financial-centre/FinTech/open-application-programming-interface-api-for-the-banking-sector/phase-approach/> and <https://www.hkma.gov.hk/eng/key-functions/international-financial-centre/open-api-for-banking-sector/phased-approach.shtml>.

*interests. What the HKMA now sets for Phase III and IV will lay a solid foundation for further opening up of banking data and functions in a prudent manner while balancing the industry needs for API development.*⁵⁴

In other words, the HKMA will publish a set of technical standards to govern the Open API framework ahead of phase 3 and 4. Since these phases will involve customer data, and the processing of transactions, the HKMA as opted for a heavier hand in guiding API standards. Thusly, the Authority will adopt a standard set of data definitions and transfer processes to ensure accurate and secure data sharing. Details on an API standard will be published in the coming months, ahead of phase 3 and 4 implementation in 2020.⁵⁵

Singapore

The Monetary Authority of Singapore (MAS) is the primary regulatory authority driving the Open Banking initiative in Singapore. Similarly to Hong Kong, the statutory objective aligned with this authority is that of Market Integrity. Yet, there is a slight yet important additional component—the promotion of innovation is an expressed regulatory priority. Specifically, MAS is mandated to ‘promote sound and progressive financial services sector through regulation and supervision.’ According to MAS, a sound and progressive financial services sector is a vital component to market integrity, providing confidence and stability across the financial system. As such, ‘MAS seeks to promote a sound and progressive financial services sector through both financial supervision and developmental initiatives.’⁵⁶

In this respect, regulatory interventions align closely with the desired development or innovation goals for the marketplace. This sentiment, perhaps, is what has made the Singaporean approach to Open Banking so widely regarded. As one of the first movers in the region, MAS first released Open Banking guidelines in 2016, ‘setting out a comprehensive framework introducing governance, implementation, use cases and design principles for APIs together with a list of over 400 recommended APIs and over 5,600 processes

⁵⁴ Mr. Howard Lee, Deputy Chief Executive of the HKMA, comments from Open API Framework for the Banking Sector: One year on. <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2019/07/20190731-3/>.

⁵⁵ <https://www.hkma.gov.hk/eng/news-and-media/press-releases/2019/07/20190731-3/>.

⁵⁶ Objectives and principles of financial sector oversight in Singapore: Monetary authority of Singapore April 2004 (revised in September 2015). <https://www.mas.gov.sg/-/media/MAS/News-and-Publications/Monographs-and-Information-Papers/Monograph--Objectives-and-Principles-of-Financial-Sector-Oversight-in-Singapore.pdf> (p. 2).

for their development.⁵⁷ By the following year, MAS, in partnership with the International Finance Corporation and ASEAN Bankers Association, launched the ASEAN FinTech Innovation Network. This network set up one of the first industry sandboxes, a ‘product-testing’ sandbox, but with a collaborative approach between industry stakeholders and a regulating authority. This sandbox effectively became a proxy for ‘an interoperable and scalable infrastructure, acting as a method to standardize banking infrastructure and data, in late 2018.’⁵⁸

Much like as in Hong Kong, the Singaporean approach to Open Banking relates primarily to an Open API framework, with minimal regulatory compulsion. That is to say, no banks are legally mandated to share data via an Open API. This is despite the fact that MAS has taken the view that wide API adoption is critical to its innovation. As such, MAS actively and with purpose encouraged financial institutions in Singapore to ‘openly develop and share APIs ahead of the license application deadlines, created resources like the Financial Industry API Register and co-developed API guidelines in partnership with the Association of Banks in Singapore (ABS).’⁵⁹

It is widely accepted that Singapore has one of the most robust API ecosystems, largely stemming from MAS’ proactive approach. Yet, as one report surmised, there does exist room for improvement and this relates primarily with a need for a legislative framework. According to a study done by Emerging Payments Association Asia (EPAA), ‘a legislative framework that lays out a common liability and complaints framework would be an important step in enabling adoption. Other measures include continuing to provide avenues for standardisation of APIs and enabling the evolution of common infrastructure and processes.’⁶⁰

Brazil

In August 2018, Brazilian policymakers published the General Data Protection Law (LGPD), which will come into effect in August 2020. Often referred to as the ‘Brazilian GDPR’, this law set out a framework for an enhanced dialogue on consumer data ownership and how personal data

⁵⁷ How DBS and OCBC global leadership shows Singapore’s success in Open Banking by FinTechnews, Singapore. 16 January 2019. <https://FinTechnews.sg/28027/openbanking/open-banking-singapore-dbs-ocbc/>.

⁵⁸ Ibid.

⁵⁹ APIs may be the key to Singapore’s Open Banking future. <https://www.pymnts.com/news/regulation/2020/apis-may-be-the-key-to-singapores-open-banking-future/>.

⁶⁰ Singapore leads APAC in Open Banking. <https://sbr.com.sg/banking-technology/in-focus/singapore-leads-apac-in-open-banking>.

is defined, and can be used, in a variety of scenarios. As this relates to Brazilian Open Banking, it legislates the legal basis for how data can be used or processed; effectively setting up the rules of engagement.⁶¹ By April 2019, the ‘Communique 33,455’ disclosed the fundamental requirements of the Brazilian Open Banking proposal.⁶² Namely, the Central Bank of Brazil (BCB) issued new guidelines aimed at the incumbent the retail banking sector, ‘with the aim of increasing competition and lowering costs for customers.’⁶³ The Brazilian Open Banking framework will fall under the supervision of the BCB, and is set to come into effect within the second half of 2020.

As outlined by the communique, the BCB’s primary objective is ‘to enhance the efficiency in credit and payments markets in Brazil by promoting a more inclusive and competitive business environment, while preserving the security of the financial system and ensuring the consumers protection.’⁶⁴ When considering the principle statutory objective, this chapter takes into account the comments by a key member from the BCB,⁶⁵ who has indicated that the primary guiding objective is most closely aligned to a Competition Mandate. This line of thinking is supported by several supporting documents, specifically the ‘Agenda BC#’ which explicitly states that ‘the Open Banking operationalization [aims] at enhancing the efficiency in credit and payments markets in Brazil through the promotion of a more inclusive and competitive business environment.’⁶⁶

Though a final regulatory approach is forthcoming, the BCB has provided us with some hints as to what is to come. In November 2019, the bank put forward draft regulation, open to public consultation. The regulatory approach is something of a ‘catch-all’, borrowing and combining elements from the UK’s Open Banking standard, the Hong Kong Open API approach, and a little bit of everything in between.⁶⁷

First, borrowing from the UK, is an emphasis on payment and current accounts, with similarly mandated guidelines and which incumbents must comply within a given time frame. Said banks (in this case the 12 largest in

⁶¹ What is the LGPD? Brazil’s version of the GDPR. <https://gdpr.eu/gdpr-vs-lgpd/>.

⁶² COMMUNIQUÉ 33,455 OF APRIL 24, 2019 Banco Central Do Brasil. https://www.bcb.gov.br/content/config/Documents/BCB_Open_Banking_Communique-April-2019.pdf.

⁶³ Why building consumer trust is key to Brazil’s Open Banking success; EY Global—14 August 2019.

⁶⁴ COMMUNIQUÉ 33,455 of April 24, 2019 Banco Central Do Brasil https://www.bcb.gov.br/content/config/Documents/BCB_Open_Banking_Communique-April-2019.pdf. (p. 1).

⁶⁵ Confidential interview conducted with representative from BCB, January 2020.

⁶⁶ https://www.bcb.gov.br/en/about/bcbhastag_competitiveness.

⁶⁷ <https://www.finextra.com/blogposting/18251/the-brazilian-way-of-open-banking-regulation>.

Brazil) will need to comply with an API standard akin to that of the UK governance structure—Data (type of), API (administration of), and Security (Technical) standards for delivering an API specification.⁶⁸ In addition to mandatory participation from major banks, the Brazilian Open Banking Framework will allow for other licensed institutions or TPPs to likewise share their data.

Much like the Hong Kong example, implementation will be divided into four phases, largely dictated by the type of data being integrated into an API. They are:

- Phase 1: Data on products, services, and channels to access banking services (by end of 2020);
- Phase 2: Data on client's personal information (address, contact details, and so on) and transactions;
- Phase 3: Services related to payment initiation transactions and credit offerings;
- Phase 4: Initiation of payments and Credit Transactions, related to currency exchange, investments, insurance, and pension funds.⁶⁹

In this respect, the scope of data types is more expansive than that of the UK's Open Banking.

Finally, one unique factor related to the BCB approach is the proposed dual regulatory structure. Though serving as the primary regulating body, setting out the scope and requirements for data sharing, the BCB may implement a 'self-regulation' approach. Though with the BCB's supervision, market players would be able to set forth a self-regulation approach, especially as related to technological standard and certain operational procedures.⁷⁰ At the time of writing this chapter, final rules on a Brazilian Open Banking framework are forthcoming, so it is unclear how this dual regulatory approach might work.

Mexico

In March 2018, Mexico published an all-encompassing 'FinTech law' (Ley para Regular las Instituciones de Tecnología Financiera), which introduced

⁶⁸ Confidential interview conducted with representative from BCB, January 2020.

⁶⁹ Adopted from <https://www.finextra.com/blogposting/18251/the-brazilian-way-of-open-banking-regulation>.

⁷⁰ Confidential interview conducted with representative from BCB, January 2020.

provisions to regulate a wide variety of financial technology models and distribution channels. This law included the first iterations of Open Banking in Mexico. When we consider which regulatory body is responsible for delivering ‘Open Banking’ this too is a bit complicated and convoluted. In order to answer this, we must first breakdown which entity will ultimately have primary scope as related to the FinTech Law. The key institutions which contribute to the policy and regulatory structures of the financial system in Mexico include the Ministry of Finance,⁷¹ Banco de Mexico⁷²; and finally the National Banking and Securities Commission (CNBV) which has an overall regulatory and supervisory role in Mexico’s financial sector.⁷³ As such, it is the CNBV which will be treated as the primary regulatory authority as related to implementing regulation prescribed within the FinTech Law.

This law is the first of its kind to formally regulate the FinTech sector, and thusly provides a regulatory framework for:

- **Electronic payment fund organisations**, responsible for the issuance, administration, redemption, and transmission of electronic payments.
- **Crowdfunding organisations**, including equity crowdfunding, P2P lending, etc.

Additionally, the FinTech Law regulates three activities:

- **Operations with virtual assets** (cryptocurrencies and related transactions)
- **A Regulatory Sandbox**
- **APIs**: with the objective of sharing financial data; either as open data, aggregated, or transactional. This type of data can be: open, aggregated, or transactional.⁷⁴

The inclusion of ‘APIs’ within this law sent Mexico to the forefront of the Open Banking conversation. This law effectively has paved the way for a Mexican Open Banking implementation, borrowing best practices from the countries that have come before (such as the UK), but also pushing the bounds and introducing the concept of Open Finance, arguably OpenBanking 2.0.

⁷¹ Effectively the policy-making arm, which plans and evaluates financial policy.

⁷² The central bank, regulating money supply, payments systems, and sets reserve requirements.

⁷³ It should be noted that a number of other regulatory bodies feed into this law, including but not limited to the Institute for the Protection of Banking Savings; the Insurance and Surety National Commission; the National Commission for the Protection and Defense of Users of Financial Services; and the National Commission for the Pension Funds System.

⁷⁴ Adopted from the FinTech Law.

In Mexico, the legal basis for regulating APIs can be found in Article 76 of FinTech Law.⁷⁵ As expressed in the ‘Pilot Project of Open Banking Standards from the UK, in Mexico’⁷⁶ Article 76 served as the main catalyst for further discussion on how best to develop a pilot project on Open Banking in Mexico. The type of data shared using these regulated APIs can broadly be described in the following three categories:

- Open financial data: information about products, branches, and ATMs, for instance. This information is usually already accessible digitally.
- Aggregated data: statistical information about operations carried out by or via regulated financial entities. This information is aggregated in such a way that personal data or data about a person’s transactions cannot be identified. Only authorised parties can access this data.
- Transactional data: information about the use of a product or service, for instance the transactions carried out, by a client belonging to a regulated entity. This information can only be shared following the client’s expressed authorisation. Payment Initiation and its related data is not included in the Mexican FinTech laws.⁷⁷

At present, the FinTech law provides some guidance and provisions as related to APIs. But it is widely acknowledged within the regulatory and supervisory ranks within the country that further considerations regarding how to regulate Open Finance exist. Fortunately, this law sets the legal foundations for such regulatory intervention.

In the first instance, it is quite likely that an Open Banking Working group will be assembled, with regulatory stakeholders from a variety of sectors. Given the role, however, of the CNBV, it is likely that any mandate given to such a group will be rooted in statutory objectives related to its base-regulator. Similar to how the CMA’s competition mandate dictates guiding principles for the OBEL, it follows that the CNBV’s Consumer Protection mandate will play a significant role in any subsequent working group.

⁷⁵ https://7da2ca8d-b80d-4593-a0ab-5272e2b9c6c5.filesusr.com/ugd/7be025_70ade1a18f3645738b7f9bc80e9fa19c.pdf.

⁷⁶ Proyecto Piloto del Estandar de Datas Bancarios Abiertos de Reino Unido en Mexico. https://7da2ca8d-b80d-4593-a0ab-5272e2b9c6c5.filesusr.com/ugd/7be025_70ade1a18f3645738b7f9bc80e9fa19c.pdf (p. 10).

⁷⁷ What is the potential for Open Banking in Mexico, recommendations and roadmap for adopting an Open Banking standard. https://docs.wixstatic.com/ugd/7be025_3c32545fddcf4436ac2a4f9fd87c9a70.pdf?__hstc=28701923.948f040f251756a8d0af8dc7c472282c.1583966970421.1583966970421.1583966970421.1&__hssc=28701923.1.1583966970423&__hsfp=2680836154&hsCtaTracking=caa33af8-c94c-4326-bea7-646e24c59e76%7Cc61c51c1-05fb-44fb-abad-eb49949dfcee (p. 19).

Under the FinTech Law, there already exists some regulatory considerations imposed upon APIs by the CNBV. In most countries, including the UK, only incumbent banks are obligated to open and share their data through the use of APIs. In Mexico, all financial entities are obliged to do so. Hence, the term ‘Open Finance’ instead of ‘Open Banking’ is employed to refer to the wider-reaching scope that includes all types of financial institutions and TPPs.

4 Concluding Thoughts

Open Banking remains an amorphous term, defined both as a regulatory intervention and as a financial technology solution. Yet, as different jurisdictions define their own unique brand of Open Banking, the concept of an Open API is being used as a synonymous (though somewhat inconsistent) term when considering Open Banking API standards.

Across all of the jurisdictions reviewed in this chapter, API standardisation is the only regulatory approach that has held true across all iterations of an Open Banking regulatory intervention. Yet standards (especially as they relate to standardising data bases, creating common technical standards and security standards) are varied, and not harmonised across the jurisdictions studied.

The UK’s Open Banking Standard is widely accepted as leader, with nearly every jurisdiction pulling from elements of the UK framework. As a regulatory intervention, the UK’s Open Banking standard has been referred to by most other jurisdictions, especially when considering API standards. In the case of Hong Kong and Singapore, for instance, Open Banking is often equated to an Open API framework, with the regulatory setting forth parameters for API standardisation, though not necessarily explicit specifications on technical standards. Conversely, the UK, Australia, Brazil, and (to a lesser extent) Mexico have dictated standards as related to the three overarching themes: data standards, specification (technical) standards, and security standards. It is important to note, no specific harmonisation of standards exist, which may limit uptake by various stakeholders as internationalisation of retail banking activities becomes more commonplace.

Throughout this chapter, I place great importance upon the Statutory Objective at the core of the respective regulator tasked with implementation within a given jurisdiction. Some similarities emerge across different jurisdictions that share regulatory objectives. In jurisdictions where a Competition mandate is the driving statutory objective, compulsory activity (especially targeted to incumbent banks) is universal. This is the case in the UK, Australia, and Brazil. Interestingly, all three of these countries are also

influenced heavily by consumer protection considerations, especially those mandated by legislation; PSD2 + GDPR in the UK, CDR in Australia, and GDPL in Brazil.

Conversely, examples of Market Integrity objectives are far less prescriptive in their Open Banking interventions, without a compulsory element. Additionally, in the case of two examples (Hong Kong and Singapore), engagement with the retail banking industry related more to market-driven partnership.

This also holds true when considering formal supervision (often in the form of a registry) of both data attribute providers, and TPPs. In the case of Hong Kong and Singapore, the role of the regulator is more akin to observer than active supervisor. In contrast, for the other jurisdictions, supervision is a critical element of their approach.

When considering the future of Open Banking, the types of data attribute providers is likely to continue expanding to go beyond banks. In that vein, Open Banking is likely to move towards 'Open Finance', with Mexico potentially spearheading this initiative. When considering the Mexican example, at present, all Financial Institutions are compelled to participate in an open data intervention. In Australia, the CDR positions data sharing to extend beyond financial markets, into telecoms and energy. In 2020, the UK's FCA will also explore the potential of an Open Finance approach.

Another limitation is that of the scope of data shared. When considering the UK's Open Banking, the OBEI, together with the FCA, has noted that the next iterations of implementation will address gaps in data types that the existing regulation missed or did not account for. As we have noted in the above pages, Australia, Hong Kong, Singapore, and Brazil have all paved the way for additional financial data, such as mortgages, cash accounts, etc. to fall under their Open Banking initiatives.

Several limitations arose when approaching this chapter. The first, and perhaps most critical, is that of a widely accepted definition on Open Banking. Given its amorphous meaning as a technical solution or regulatory regime, it is difficult to accurately describe the different iterations of Open Banking effectively. This difficulty is compounded by the fact that different jurisdictions are defining the parameters of Open Banking in different ways.

Given the nascence of the topic, significant holes in the literature exist. Primary source material is mostly related to published regulatory papers, consultations, laws, or memorandums. In many jurisdictions, regulators do not update or keep appropriate corresponding articles live, and language issues prevented me from definitively capturing nuances in certain jurisdictions.

Figure 2 outlines each of the case-study jurisdictions by the primary regulating authority and statutory objective influencing its specific approach to Open Banking. As one will note, the figure highlights the different components a regulator may employ when pursuing their specific approach to Open Banking. This is a non-exhaustive list but includes the tendencies that are re-occurring across a variety of jurisdictions. Finally, this figure reviews the types of Data and associated activities included within an Open Banking initiative. As such, this figure allows for a comparative review of the different case-studied regulators by approach and data attributes, against their statutory objective.

Jurisdiction	Regulatory/Supervisory Authority	Primary Statutory Objective	Regulatory Approach Check-List					Types of Data and Associated Activities						
			API Standardisation	Forward Regulatory Approach	Non-prescriptive Governance Framework/Agreement (or related instrument, contract, consumer protection)	Current Use of Data	Primary Information	Additional financial products (i.e. mortgages, pensions, insurance or cash accounts)	Secondary Data (including third party)	Secondary Information (if applicable)	Primary Statutory Objective			
United Kingdom	Competition and Markets Authority	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
France	Authorite de Competition	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Australia	Competition and Consumer Commission	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Hong Kong	Competition Commission	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Hong Kong	Monetary Authority of Hong Kong	Monetary Authority of Hong Kong	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Singapore	Competition Commission	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Brazil	Nacional Bank and Securities Commission	Consumer Protection	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Malawi	Competition Commission	Competition	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

* Influenced heavily by Consumer Protection obligations

** Y or related to security concerns in para 3 and 4

*** MAS maintains a Financial Industry API Register which tracks API by functional category

**** M initial under Force of Law, further work forthcoming

Fig. 2 Case-study jurisdictions by primary regulating authority and statutory objective

Welfare Implications



Technological Innovation and Discrimination in Household Finance

Adair Morse and Karen Pence

1 Introduction

Technology has the potential to be a tremendous anti-discriminatory force in household finance. Technology can displace human discretion in decision-making—a traditional source of bias and discrimination—with more objective decisions based on data and algorithms. Technology can bring access to financial services to households who were previously excluded from these markets because of financial or geographic isolation; this exclusion itself can stem from discrimination. Technology can reduce the costs associated with extending credit and make it profitable to extend credit to households out of the mainstream majority community. Technology can increase access to information about the availability and pricing of financial products and services and thereby level the information playing field.

However, a positive outcome for the role of technology in mitigating discrimination is not pre-ordained. Technology inherently has no animus, but is not immune from being discriminatory. Some households have more access

A. Morse (✉)

Haas School of Business, University of California, Berkeley, CA, USA
e-mail: morse@haas.berkeley.edu

U.S. Department of the Treasury, Washington, DC, USA

K. Pence

U.S. Department of the Treasury, Federal Reserve Board, Washington, DC, USA

to and more facility with technology than others, such that even innovations may reinforce patterns of exclusion. Technology can allow firms to target advertising and product offers very precisely to consumers, raising the possibility that households have different information sets and even face different prices, sometimes in breach of fair lending, equality, public accommodation, and civil rights laws. Decision-making by financial services providers via algorithms incorporates thousands of variables and presents courts, policymakers, and regulators with complex questions about how to think about and detect discrimination. And although algorithms have no inherent bias, they can incorporate the biases embedded in the broader culture through the datasets used in their development and through the biases of their development teams.

This chapter brings together these benefits and friction points of how technology in finance can affect discrimination. The evidence indicates that technology is a powerful force for reducing discrimination stemming from human discretion (“taste-based” discrimination in economics parlance). But the net effect of technology as an abater of discrimination, especially looking into the future, is not obvious, and depends heavily on resolutions to legal and regulatory uncertainties surrounding the use of algorithms in what economists call “statistical” discrimination. If whether technology is net positive or negative for discrimination is the thread woven through our chapter, the overall tapestry of our contribution lies in identifying the technological implementations that can lead to discrimination, particularly focusing on the interactions between financial service providers and households in human discretion, algorithmic decision-making, and innovation and inclusion. We identify five such gateways for discrimination: (i) human involvement in designing and coding algorithms, (ii) biases embedded in training datasets, (iii) practices of scoring customers for creditworthiness based on variables that proxy for membership in a protected class, especially through digital footprint and mobile data, (iv) practices of statistical discrimination for profiling shopping behavior, and (v) practices of technology-facilitated advertising, including ad targeting and ad delivery.

Within these implementations of technology, we further identify four regulatory “frontlines.” We use the term *frontline* to connote two sentiments—a situation of uncertainty (in particular, as to whether the regulatory status quo will remain) and a setting of potential conflict (as legal protections of individuals confront forces of business use of technology). How these legal and regulatory frontlines are resolved will affect whether technology is on net positive or negative in the long run for discrimination.

Our regulation frontlines concern regulatory uncertainty concerning: (i) whether a variable is “correlated enough” with a protected class to be discriminatory itself, (ii) the use of input-based enforcement of large dataset algorithms versus output-based compliance, (iii) the extent to which privacy laws restrict algorithmic provision of financial services, and (iv) the applicability of public accommodation laws (also called equality laws) to disparities in access to online and mobile provision of financial services. These points of uncertainty affect not just legal tensions of how regulators and courts will act, but also the ability of financial service providers to innovate and the incidence of the benefits of innovation to consumers.

Our chapter builds heavily on the works of other scholars that examine various settings or specific aspects of discrimination. We highlight these works as we proceed. Our contribution is in the amalgamation and analysis of ideas toward understanding the gateways of discrimination entering technological finance, the frontlines of regulation, and the weights for and against technology as an abater of discrimination in financial services. In the process, we gain the insight of just how drastically technology has changed the way discrimination manifests itself in financial services.

2 Views of Discrimination: Lawyers and Economists

The U.S. has comprehensive federal laws prohibiting discrimination in lending, and a patchwork of state and federal laws that cover, less comprehensively, other financial services.¹ Evans (2017) provides a review of these laws with an emphasis on their implications for fintech firms. In the United Kingdom (UK) and Canada, discrimination in financial services is prohibited under broader antidiscrimination laws.² The European Union (EU) recognizes non-discrimination as a fundamental right, but relegates specific legislation to member states, who in turn vary in their attention to discrimination legislation and enforcement. In practice, financial service providers have had more freedom in continental Europe to use protected characteristics for profiling, but this is changing, as the EU and UK take a leadership role in regulating the use of technology in finance.

¹ The main fair lending laws in the United States are the Equal Credit Opportunity Act and the Fair Housing Act.

² The Equality Act 2010 prohibits discrimination in the provision of services in the UK. See Hale (2018) for a discussion of equality under the law in the UK. The Canadian Human Rights Act prohibits discrimination at the federal level in Canada. Discrimination may also be regulated at the provincial level for financial service providers that operate in only one Canadian province or territory.

Discrimination laws cover a varying set of protected classes. Individuals are usually safeguarded against discrimination based on race, ethnicity, religion, marital or family status, and disability, and sometimes on additional characteristics such as age, gender, national origin, sexual orientation, gender reassignment, political views, genetic or biometric information, veteran status, and use of social safety nets.

Discrimination laws in principle cover all the steps and practices involved in offering a financial service. However, the enforcement by which this principle is carried out varies from country to country on at least two dimensions. First, countries may differ in the intensity of the focus on the steps. The U.S. antidiscrimination laws and their implementing regulations generally have greater specificity as to the steps and practices covered; in fair lending, for instance, a lender in the U.S. cannot discriminate in advertising, credit risk assessment, and pricing of a loan. The UK's broad-based laws have some individual requirements and carve-outs but in general have less specificity than the U.S.³ Second, countries may differ in the specificity of sectors to which antidiscrimination laws apply. Again, the U.S. code is more specifically written, delineating housing, credit, and employment as sectors with particularly detailed regulations. An advantage to specificity in preventing and enforcing is the attention to the particulars for the steps and sectors listed in the laws. An advantage to generality is in flexibility to expand to considering new steps and sectors as the provision of financial services expands and changes with technology. Of course, the question in the more general case—a question being played out across the different country jurisdictions in the EU—is whether a country will delve into discrimination compliance within the steps of provision if the law does not explicitly say so.

Finally, all discrimination laws speak to direct discrimination—treating individuals differently on the basis of protected characteristics such as race, ethnicity, or gender. This is called *disparate treatment* in the U.S. U.S. regulators distinguish between “overt evidence” of disparate treatment, when a lender openly discriminates on a prohibited basis, and “comparative evidence” of disparate treatment, when a lender treats an applicant differently based on a prohibited basis.⁴ Comparative evidence can encompass treating individuals differently on the basis of variables that are highly correlated with a prohibited basis. In the U.S., variables such as grey hair (for age in employment decisions) or zip codes (for minority neighborhoods in credit decisions) fall under

³ This may be changing, as a new regulation took effect in June 2019 that forbids any advertising that includes gender stereotypes that are likely to cause harm (Safronova, 2019).

⁴ FDIC Consumer Compliance Examination Manual, Fair Lending Laws and Regulations, IV—1.1–1.2, September 2015, <https://www.fdic.gov/regulations/compliance/manual/4/iv-1.1.pdf>.

this category and are generally considered direct discrimination, since these variables are masked versions of the original prohibited basis.⁵ Some discrimination laws also encompass indirect discrimination, which is when a policy or practice that on its face seems neutral disadvantages a protected group indirectly through other variables. The Australian Human Rights Commission provides the example of a public building being only accessible by stairs as representing indirect discrimination against people with disabilities. This example is *indirect* because the lack of accessibility of a building is presumably not done to preclude people with disabilities but rather to save on costs.⁶ In the U.S., the term *disparate impact* maps roughly to indirect discrimination.

In contrast to legal views of discrimination, economists view discrimination through the lens of whether it is “taste-based” or “statistical.” Under taste-based discrimination (Becker, 1957), decision-makers get utility from engaging in prejudice, and are willing to sacrifice other priorities—such as hiring the most productive workers possible—in order to satisfy their biases. The much-cited culmination of Becker’s theory is that taste-based discrimination cannot persist because other employers, who do not have a taste for prejudice, will hire workers based solely on their productivity. These non-discriminating firms will be more profitable than their prejudiced competitors, and the prejudiced firms will go out of business. As we discuss in section III, this culmination may not play out in practice.

Under statistical discrimination (Arrow, 1973; Phelps, 1972), discrimination results from the practice of using variables as statistical discriminants to uncover unobserved variables. There are two crucial differences between statistical and taste-based discrimination for our purposes.

First, statistical discrimination does not require employers or other decision-makers to have animus or negative taste toward a protected category (non-whites in the Becker formulation). Rather, decision-makers engage in statistical discrimination because they are missing information on a characteristic that is key to their decision, such as credit risk in the case of lending. In the formative theory models, a lender that lacks such information may try to recover proxies for credit risk by using the average credit risk of a group, where the group is defined by gender, race, ethnicity, or other characteristic. In practice, applying the averages by protected groups is illegal, but lenders may use other variables that correlate with a protected category to recover credit risk, and thereby implement statistical discrimination in a more general way than the original theories.

⁵ Rothstein (2017) discusses the historical roots of redlining. Pop (2013) describes a similar debate that unfolded in Germany.

⁶ <https://www.humanrights.gov.au/quick-guide/12049>.

Second, statistical discrimination, unlike taste-based discrimination, is profit-maximizing for the financial service providers. This finding implies that the target for using statistical determinants from the firm (and economist) perspective is profits, not uncovering the unobserved component of credit risk. We discuss in section *IVa* how this economists' concept of statistical discrimination sometimes misaligns with the legal view. As a preview here, we note that the possibility that financial service providers could be illegally discriminating while profit-maximizing is an uncomfortable juxtaposition of the economist's view of discrimination with the law.

3 Human Decisions and Discretion

3.1 Discriminatory Discretion Ameliorated by Technology

Historically, lenders exhibit patterns in providing financial services that appear consistent with taste-based discrimination against certain types of individuals, even when acting on these biases has resulted in lower profits (Charles and Hurst, 2002; Bayer, Ferreira, and Ross, 2017; Alesina, Lotti, and Mistrulli, 2013; Deku, Kara, and Molyneux 2016; Dobbie, Liberman, Paravisini, and Pathania, (forthcoming); and Bartlett, Morse, Stanton, and Wallace, 2019a). In Becker's theory, taste-based discrimination is competed away by market forces. However, if the market is not fully competitive, or if the foregone profits associated with employees who discriminate are fairly small, taste-based discrimination can persist. This type of discrimination is particularly likely to emerge in settings where decision-makers have discretion.

Technology has the potential to limit discretionary discrimination by providing information about financial services more broadly and at lower cost and by limiting the face-to-face interactions that appear to facilitate discrimination. Scott Morton, Zettelmeyer, and Silva-Russo (2003), for example, found that Black and Latinx car purchasers paid more than white purchasers when the sales negotiation took place in person but not when it occurred on the Internet.

When humans are removed fully from negotiations, the decision-making becomes algorithmic, which has been found to reduce costly discriminatory discretion in many settings. For instance, Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2018) show that a machine learning algorithm outperforms human judges in predicting which defendants will skip their

next court appearance or commit crimes while out on bail, and does so without increasing racial disparities in the probability of being released on bail. In the realm of lending, Dobbie, Liberman, Paravisini, and Pathania (forthcoming) show that a high-cost lender in the UK would increase profits and reduce bias if it used a machine learning-based algorithm to make lending decisions instead of relying on the judgment of lending examiners. Bartlett et al. (2019a) find that the discrepancy between the rates charged to white and Black/Latinx borrowers is lower for algorithmic lenders than conventional lenders, and that algorithmic lenders show no disparities in mortgage rejection rates.

3.2 Discriminatory Discretion Enabled by Technology

Yet technology does not always remove discretion, because humans remain involved in technology processes—either through the role of peer input in platforms or through coding. The new modes of human involvement in technology-provided financial services shift some of the fault lines of discrimination from banking and loan officer discretion to coding and the crowd.

Peer-to-peer platforms provide a setting where humans remain involved in a technology-induced market. Just as in brick-and-mortar lending, when online platforms are information portals for human decision-making (by the peer investors) and include pictures and names, discretion biases decisions toward taste preferences and in-group biases. For instance, Edelman, Luca, and Sverisky (2017) find that AirBnB applicants with distinctively African-American names are less likely to be approved. In the lending context, Ravina (2019) studies the platform prosper.com's use of the photos of borrowers. She finds that attractive borrowers are more likely to get a loan and more likely to default. Online platforms have little incentive to rein in this behavior because providing pictures and names appear to give consumers more trust in participating in transactions with strangers (Edelman and Luca, 2014, Doleac and Stein, 2013) and because the Communications Decency Act may shield the providers from liability (Edelman and Stemler, 2018).

Coding and data inputs are another avenue through which human discretion can lead to discrimination in technological finance. Algorithms are ultimately designed by humans, even if the mechanics are handled by artificial intelligence. How programmers set up the optimization problem, classify the data, and choose the training data sets can influence the output of the algorithm (Barocas and Selbst, 2016). In 2017, for example, Amazon scrapped an artificial intelligence tool for reviewing resumes because it systematically

discriminated against women. The tool was trained on the resumes of past job applicants, most of which were men, and so it learned to penalize words such as “women’s” (Dastin, 2018). As is seen in this example, training data sets can facilitate discrimination either because they embed the existing prejudices in society or because they are not representative of the broader population. The stark under-representation of Blacks, Latinx, and women among programmers means that algorithms are unlikely to be designed by diverse teams that might spot some of these issues (Turner, 2018).

Discrimination Gateway #1: Human involvement in designing and coding algorithms can lead to discrimination.

Caliskan, Bryson, and Narayanan (2017) provide a systematic study of the prejudice that can result from biased data. They find that word-association algorithms trained on text from the Internet were more likely to associate European-American names with pleasant attributes such as “health” and “honest” and African-American names with unpleasant attributes such as “poverty” and “failure.” The algorithms were also more likely to associate male names with words about careers and female names with words about families, and to associate certain types of occupations with gender. They note, for example, that at the time of their paper, Google translated “O bir doktor. O bir hemşire.” as “He is a doctor. She is a nurse.” despite the fact that Turkish pronoun *O* has no gender.

Other algorithms may perform poorly because they are trained on non-representative data. Buolomwini and Gehr (2018) show that facial recognition software, which is developed using machine learning techniques, is less accurate in classifying the gender of females and of dark-skinned individuals. The three classifiers that they examined had error rates of less than 1% for light-skinned male faces and from 20 to 34% for dark-skinned female faces. The authors note that darker-skinned faces appear to be less represented in some datasets that are used to train the software. They also cite the finding of Roth (2009) that default camera settings are set to expose lighter-skinned faces. Similarly, Blodgett and O’Connor (2017) show that language-identification software is less accurate in identifying as English shorter-length Tweets written in the dialect of American English referred to in their paper as African-American English.

Discrimination Gateway #2: Biases can be embedded in training datasets.

4 Algorithmic Decisions

By reducing the role of human discretion, algorithmic decision-making has the potential to decrease discrimination in financial services. But algorithmic decision-making may also amplify discrimination because the opacity and complexity of the algorithms may mask the fact that the use of variables can cause disparities in outcomes against protected groups that is illegal under some antidiscrimination laws. In this section, we focus on three interactions between households and financial service providers—credit risk assessment, pricing of financial services, and advertising—where the introduction of new algorithmic modeling techniques may cause disparities not justified in antidiscrimination laws.

4.1 Statistical Discrimination

In this section, we bridge the economists' concept of statistical discrimination with the legal framework. Statistical discrimination, in the eyes of economists, is a solution to a signal extraction problem, where the signal provides data on a fundamental skill or attribute that is otherwise unobservable. In the eyes of the law, businesses are allowed to use proxies for these unobserved factors if it can be justified as a legitimate business necessity and if the use of these proxy variables does not have a disproportionate effect on individuals in a protected category.

Statistical discrimination can help firms model unobserved variables that are key to their decision-making. For example, in the context of credit risk scoring, recent evidence in Pope and Sydnor (2011) and Ravina (2019) demonstrates that unobserved fundamental credit risk is correlated with race even controlling for modern-day, sophisticated observable measures of credit risk. This finding suggests that if proxy variables correlated with hidden credit risk fundamentals were available, statistical discrimination might enable lenders to do a better job in modeling credit risk.

Yet depending on the legal setting, profiling individuals, even without taste-based intent to discriminate, with proxy variables that results in disparate outcomes by protected categories is illegal under two conditions and may be ruled in breach (depending on jurisdiction) in another three situations.⁷ We refer to these as the *illegal practices* and *frontiers* in the discussion below.

⁷ We use the term statistical discrimination for these activities or outcomes, but some prefer to use this term only for legal use of variables to profile individuals according to the law.

4.1.1 Practices Aimed at Statistical Discrimination that are Illegal

First, the use of a protected category (in most jurisdictions) or a variable highly correlated with the protected category (in many jurisdictions) as a statistical determinant is illegal, even if the variable is also correlated with the unobservable target of the statistical discrimination application. (The term “target” refers to the business necessity motive—creditworthiness in finance, flight risk in bail setting, or productivity in labor decisions.) This rule is clear enough on the use of the protected category variable, but the threshold of what is “highly correlated” is ambiguous and defined in a context-specific application by courts or regulators. For example, in February 2019, the New York City Commission on Human Rights issued a prohibition on businesses using hairstyles to sort individuals for decisions, because of a high correlation with race (Stowe, 2019). Furthermore, the concept of what is highly correlated becomes increasingly complex as the number of variables used in an algorithm expands with technology. Because of this ongoing tension, we present this concern as our first regulatory uncertainty frontline:

Regulatory Frontline #1: Uncertainty exists in the extent to which courts, regulators, and policy makers will expand the list of highly correlated variables for exclusion in statistical discrimination in an era where it is easier both for financial service providers to amass such variables and for compliance officers and plaintiffs to calculate statistically significant correlations.

Second, practices aimed at statistical discrimination are not legal in the U.S. if the target motivating the statistical discrimination is not defined as a court-justified motive. For lending, U.S. courts have been explicit in ruling that the target is credit risk assessment and that profit motives beyond credit risk are not legal reasons for statistical discrimination. In particular, Bartlett et al. (2019a) highlight three cases stating this precedent:

In *A.B. & S. Auto Service, Inc. v. South Shore Bank of Chicago*, 962 F. Supp. 1056 (N.D. Ill. 1997), the ruling reads: “...the defendant-lender must demonstrate that any policy, procedure, or practice has a manifest relationship to the creditworthiness of the applicant...”. This language again appears in *Lewis v. ACB Business Services, Inc.*, 135 F.3d 389, 406 (6th Cir. 1998): “The [ECOA] was only intended to prohibit credit determinations based on ‘characteristics unrelated to creditworthiness,’” and in *Miller v. Countrywide Bank, NA*, 571 F.Supp.2d 251, 258 (D. Mass 2008): “[rejecting arguments] “based on subjective criteria beyond creditworthiness.”

These directives by the court matter because lenders might rationally price loans strategically to maximize profits, not just to profile individuals on credit risk. Using profits as the target may not be legitimate under the law if lenders take higher profit margins above costs from a protected category group. We return to this point when talking about pricing of financial services.

In this U.S. setting that focuses on the business necessity motive for statistical discrimination, the determination by courts, regulators, and compliance officers that certain practices are illegal is based on determining whether a specific input, such as a variable, dataset, or process, has led to a disparity in outcomes against a protected group. Recently, however, researchers have begun to advocate for approaches other than this traditional approach of simply analyzing an algorithm's decision process through consideration of its inputs. Pope and Sydnor (2011) offer an approach that is based on constructing *corrections* to disparities created against protected groups by input variables, as opposed to actions to exclude such inputs.⁸ Cowgill and Tucker (2020) go further, arguing that “regulations focusing on *outcomes* exhibit more flexibility, fewer loopholes, greater efficiency and stronger incentives for innovation.” Yet Bartlett, Morse, Stanton, and Wallace (2019a) contend that U.S. law demands an approach that first considers the justifiability of business necessity criteria in the input use, which means an output-assessment defense to a discrimination claim will not hold up in court. In a similar vein, Kleinberg, Ludwig, Mullainathan, and Sunstein (2018) argue that antidiscrimination prevention and enforcement should focus on the training procedure⁹ with an emphasis on the transparency of all code and datasets used to determine how inputs are used in algorithmic screening.¹⁰

Regulatory Frontline #2: Algorithmic decision makers are increasingly advocating for the use of output-testing or correction approaches in compliance and court defenses, creating a fault line tension as to whether input-based compliance, which the law demands, can be challenged.

⁸ See Altenburger and Ho (2018) for implementation frictions with this method.

⁹ A training procedure in credit screening uses a dataset of inputs, lending decisions, and loan success outcomes (repayment and profits) to determine the algorithm's parameters for screening new applicants.

¹⁰ Further interesting reading on the application of statistical discrimination is in Bohren, Haggag, Imas, and Pope (2019), who provide a context for how behaviors relegate statistical discrimination practices into accurate and inaccurate sortings.

4.1.2 Frontiers of the Legality of Statistical Discrimination

In Europe, the landmark European General Data Protection Regulation (GDPR), which took effect in May 2018, creates challenges to the use of direct and indirect variables in profiling individuals for financial services. The GDPR applies to any company that processes the personal data of EU residents, regardless of where that processing takes place. The law forbids processing “special-category” data (race, ethnicity, sexual orientation, religion, genetics, biometrics, etc.) unless the individual has given consent or such data collection is in the public interest. This provision makes direct use of these special-category variables illegal in all countries in the EU without consent. (An interesting question is whether the probability of consenting to such data collection itself varies systematically across groups.) The law also limits the ability of firms to use only an algorithm to make decisions that affect humans significantly, such as job offers or extensions of credit, without explicit consent. Firms must disclose to consumers the types of data and the logic used by the algorithm. Further, if a consumer is denied credit or other benefits on the basis of an algorithmic decision, firms must provide consumers the option to appeal the decision to a human decision-maker. It is yet to be seen whether individuals’ power over consent has leverage in preventing discrimination. Goodman and Flaxman (2017) discuss the frictions of the implementation of the GDPR for algorithmic decision-making and machine learning, and comment on how discrimination cannot be eliminated by excluding protected categories data because firms can still engage in statistical discrimination via variables that are correlated with protected classes.

Regulatory Frontline #3: Considerable uncertainty surrounds the eventual scope and implementation of the GDPR, especially as it is interpreted to encompass statistical discrimination on variables correlated with special-category data.

The second frontier challenge is being played out in the U.S., where technology has brought out new interest in the language about “unfair or deceptive acts or practices” regarding data from the Federal Trade Commission Act and the Dodd-Frank Wall Street Reform and Consumer Protection Act. A number of compliance risks associated with these laws in the context of data inputs and algorithmic decision-making are found in Evans (2017) and Federal Trade Commission (2016), including discussions of recent actions against companies for failing to disclose that borrowers’ credit limits could be reduced based on a behavioral scoring model, for misrepresenting how data

collected online was used, and for selling data to customers that a company had reason to know would use the data for fraudulent purposes.

The third frontier challenge is playing out concerning what are called public accommodation laws (in the U.S.) or equality laws (in Europe, the United Kingdom, and Canada). These statutes govern the equal rights of individuals to access business establishments and services. In the U.S., these laws are primarily at the state level and vary in the types of individuals that are protected and in how comprehensive the laws are.¹¹ For example, the Unruh Act in California states that

All persons within the jurisdiction of this state are free and equal, and no matter what their sex, race, color, religion, ancestry, national origin, disability, medical condition, genetic information, marital status, sexual orientation, citizenship, primary language, or immigration status are entitled to the full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever.

California Civil Code § 51(b).

The use of these state laws to police discrimination in online platforms is new and is rapidly evolving. Initially, legal actions against online service providers under the Unruh Act focused primarily on whether websites were accessible to people with disabilities.¹² In 2019, the California Supreme Court established that the Unruh Act governs discrimination against people who intend to transact with websites, regardless of whether they actually do so.¹³ The most recent focus is discrimination in advertising. In July 2019, Governor Cuomo asked the New York Department of Financial Services to look into Facebook's policy of allowing advertisers to target ads by protected classes, and in November 2019, a lawsuit filed in California alleged that this same Facebook policy violates the Unruh Act (*Opiotennione v. Facebook Inc.*, hereafter *Opiotennione*).¹⁴ These developments mark a possible expansion of the scope of public accommodation laws to discrimination by algorithms. It also marks

¹¹ According to the National Conference of State Legislatures: "Five states... do not have a public accommodation law for nondisabled individuals. All states with a public accommodation law prohibit discrimination on the grounds of race, gender, ancestry and religion. In addition, 18 jurisdictions prohibit discrimination based on marital status, 25 prohibit discrimination based on sexual orientation and 21 prohibit discrimination based on gender identity. Nineteen jurisdictions also prohibit age-based discrimination in areas of public accommodation." <https://www.ncsl.org/research/civil-and-criminal-justice/state-public-accommodation-laws.aspx>.

¹² *Thurston v. Midvale Corp.*, 39 Cal. App 5th 634 – Cal: Court of Appeal, 2nd Appellate Dist., 8th Div. 2019.

¹³ *White v. Square*, 7 Cal. 5th 1019, 1025 (2019).

¹⁴ Complaint, *Opiotennione v. Facebook Inc.*, No. 3:19-cv-07185 (N.D. Cal. filed Oct. 31, 2019).

an expansion of the venues available for redress in the U.S. for discrimination in financial services other than lending. As noted earlier, federal discrimination regulations are more comprehensive for lending than for other financial services.

Regulatory Frontline #4: It is yet to be seen as to whether jurisdictions will rule that public accommodation and equality laws apply to the policing of algorithmic information provision and access to mobile and online financial services.

4.2 Credit Risk

Lenders may have a profit-maximizing motive to discriminate in loan underwriting inasmuch as membership in a protected class could be correlated with a credit risk factor that is difficult for the lender to observe, even after conditioning on observable measures of credit risk. Such unobserved risk factors are likely themselves inextricably tied with past or present discrimination. For example, because of long-standing discrimination in the labor market or in the education system, family resources (an often hidden fundamental in credit risk) might be lower for some groups than others. We emphasize that the possible existence of such a profit-maximizing motive does not mean that such profiling is legal.

The advent of technology has changed this dynamic in two ways. First, data advances have made it easier to measure underlying creditworthiness and thus reduced lenders' need to rely on flawed proxies. For example, lenders now extract information from transaction-account data to create measures of expected cash flow.¹⁵

Second, technology has amplified the practice of using correlated variables as proxies for missing fundamentals on credit risk, sometimes resulting in statistical discrimination against protected classes of households. What distinguishes credit risk assessment today from a long history of lenders using proxies is the sheer number of variables involved and the complexity and opacity of the algorithms. The use of new variables is not necessarily problematic. Upstart Network, for example, asserted in a submission with the Consumer Financial Protection Bureau that its use of education and employment history as underwriting variables expanded access to credit without displaying any disparities that required further fair lending analysis (Ficklin and Watkins, 2019).

¹⁵ One evaluation of these cash-flow measures found that their use in underwriting appears to expand access to credit without creating a disparate impact on protected groups (FinRegLab, 2019).

This is not always the case, but the evidence is yet thin. What we do know comes from implementations in countries with fewer restrictions on the use of proxy variables as it relates to discrimination. For example, Berg, Burg, Gombović, and Puri (2020) showed that “digital footprints” such as the type of device (tablet, computer, phone), operating system (Windows, iOS, Android), and email provider predicted default rates among the customers of a German lender. It is possible that credit risk assessment based on digital footprints would fail discrimination suits, as the case could be made that such variables disproportionately affect a protected group beyond any effect that operates through creditworthiness. Likewise, BJORKEGREN and GRISSEN (2019) generated measures of creditworthiness from mobile phone usage data in a South American country and showed that these measures predicted default. Again, mobile phone usage might be additionally correlated with a protected category, even beyond its correlation with credit risk fundamentals.

Discrimination Gateway #3: Practices of scoring customers for creditworthiness based on variables (including the now-pervasive digital footprint and mobile data) that correlate with membership in a protected class are at risk to lead to discrimination.

4.3 Pricing

Pricing is fertile ground for discrimination. There is wide variation in the amount that individuals pay for financial services, even for the same loan originated in the same market on the same day (Bhutta, Fuster, and Hizmo, 2019; Alexandrov and Koulayev, 2018). Within that variation, women and Black/Latinx men appear to pay higher rates on loans than white men (Woodward and Hall, 2012 (and references therein in online appendix A), Bartlett et al., 2019a). Class-action suits and enforcement actions brought against mortgage lenders in the aftermath of the financial crisis provided considerable evidence of discretionary pricing policies that discriminated against Black and Latinx borrowers (Ayres, Klein, and West, 2017).

This pricing dispersion persists, in part, because many individuals do little comparison-shopping for financial services products and because some markets are more competitive than others. Alexandrov and Koulayev (2018) noted that about half of mortgage borrowers do not shop before taking out a mortgage. Bhutta, Fuster, and Hizmo (2019) find that paying a lower mortgage rate is associated with whether the borrower considered multiple lenders or consulted mortgage websites. In their paper, this shopping behavior partly explained why borrowers with lower incomes, wealth, and credit scores

paid higher mortgage rates. Although these particular characteristics are not protected classes, Woodward and Hall (2012) show that limited shopping is one of the reasons why Black and Latinx borrowers pay higher mortgage broker charges.¹⁶

Bartlett et al. (2019a) provide supporting evidence that even in a setting in which lenders bear no credit risk, Black and Latinx borrowers pay higher prices for mortgages issued by both traditional and FinTech lenders. Their contribution does not provide direct evidence tying rates to shopping behavior as in Woodward and Hall (2012). However, because Bartlett et al. can rule out omitted credit risk variables that could confound the interpretation of other studies, the disparate pricing result must come from either profiling for shopping or the competition environment, and thus is a complement to Woodward and Hall (2012).

Indeed, algorithms can infer the propensity of an individual to shop around, as well as other individual-specific factors that affect a consumer's willingness to pay for a product. Thus technology allows firms to make these inferences with far greater precision than was available previously and to use that information in a pricing strategy. Donnelly, Ruiz, Blei, and Athey (2019) (also see Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, forthcoming) present a machine learning model, for example, that identifies which consumers are most price sensitive in their demand for a given product and allows for personally targeted price discounts. While this behavior is profit-maximizing for the firm, U.S. courts have ruled that profit motives beyond credit risk are not legal reasons for statistical discrimination in pricing, as discussed in section IV.a. and in Bartlett et al. (2020b).

We also note that the economist concept of “profit-maximizing for the firm” can be observationally equivalent to the consumer advocate's concept of “profiteering from vulnerable groups.” The connection between pricing and shopping behavior is complicated, for example, because shopping behavior also reflects differences in search costs across groups. Ayres (1991) describes some reasons why women and nonwhite men might face higher search costs in the auto retail market. Shopping is also linked to advertising, and as we discuss in the next section, there is some evidence that Black and Latinx individuals are less likely to be shown ads that provide them with useful information. While firms may not have much control over shopping behavior per se, they do have control over the information that consumers receive and the

¹⁶ Tabulations from the 2016 National Survey of Mortgage Originations, however, suggest that Hispanics and nonwhites are a bit more likely than non-Hispanic whites to consider more than one lender (Avery, Bilinski, Clement, Critchfield, Frumkin, Keith, Mohamed, Pafenberg, Patrabanish, and Schultz, 2018, Table 10).

way choices are framed. We turn next to a discussion of how technology has changed that dynamic.

Discrimination Gateway #4: Practices of statistical discrimination that profile shopping behavior can lead to discrimination, even if the firm's statistical discrimination is motivated purely by the economics of profit maximization.

4.4 Advertising and Information

Companies always have had an incentive to tailor their advertising content, targeting, and delivery to different markets. In this section we discuss how technology is changing this playing field and how these changes interact with the existing laws on discrimination. Evans and Miller (2019) provide a fuller treatment of the fair lending implications of these issues.

Evidence from mortgage advertising underscores the existing differences in information delivered to target audiences, even before technology. Perry and Motley (2009) show that prime borrowers were more likely to be shown advertisements with detailed information that helped them make better financial decisions, whereas subprime borrowers were more likely to be shown information that played on their fears. Some of the ads in Perry and Motley (2009) were explicitly aimed at minority borrowers.¹⁷ This Perry and Motley finding is consistent with a lawsuit filed by the Attorney General of Illinois that alleged that Countrywide Financial Corporation aggressively marketed subprime mortgages to Black and Latinx borrowers.¹⁸

Although it is unclear whether technology affects firms' incentives to provide helpful or deceptive information, particularly in ways that differ by protected class, technology clearly amplifies the ability of firms to target their

¹⁷ Perry and Motley (2009) documented the differences between prime and subprime mortgage advertisements in a sample of ads placed on television, radio, and print from 2005 to 2007. They found that 50% of ads targeted to prime borrowers contained specific loan details, compared with 21% of ads targeted to subprime borrowers. Mortgage ads that positioned the lender as a solution for borrowers "drowning in debt" were placed almost exclusively in publications targeting Black and Latinx households. These ads contained very little information on the actual loan terms and thereby increased the borrower's dependence on the mortgage lender.

¹⁸ Complaint, *People of the State of Ill. v. Countrywide Fin. Corp.*, No. 10-ch-27929 (Cir. Ct. of Cook Cty., Ill. filed June 29, 2010). Another example of the discrepancy between prime and subprime advertising campaigns can be seen by contrasting the findings of Grundl and Kim (2019) and Gurun, Matvos, and Seru (2016). The first paper showed that the direct-mail mortgage refinance advertisements to prime borrowers appeared to improve their welfare by prodding them to refinance when it was in their interest. The second paper looked at subprime mortgage lenders and found that the lenders that advertised more heavily in print media and direct mail also tended to provide less favorable terms to borrowers. This pattern was concentrated in zip codes with more minority, low-educated, or low-income residents.

advertising. As Athey and Luca (2019) note, “Older media, such as print and television, do not allow for showing different advertisements or tracking behavior at the individual consumer level” (p. 219). In contrast, in Internet and mobile advertising, advertisers create an advertisement and specify their target audience. Often multiple advertisers are looking for the same audiences, and an algorithmic auction ensues based on the relevance of the ad and the advertisers’ relative willingness-to-pay.¹⁹ All of these mechanisms result in systematic differences in the ads that different demographic groups see.

In online marketing, “ad targeting” is the term that applies to when advertisers choose to have their ads shown only to certain groups. For ads that concern housing, employment, or credit, ad targeting to a protected class is illegal under the same federal laws that govern discrimination in other aspects of these markets.²⁰ Tech companies are increasingly being held accountable for facilitating such targeting.²¹ Open questions are threefold. First, is this targeting legal for advertisements for financial services other than credit? The current test case, mentioned in section IV.*a.*, concerns whether the court will decide that public accommodations law prohibit discrimination in advertising for financial services other than credit and has been brought under California’s Unruh Act. *Opiotennione* alleges that Facebook denied older and female users “the opportunity to learn about and obtain financial services...over the past three years (or longer) due to Facebook’s discriminatory advertising and business practices and its aiding and abetting of financial services companies’ discriminatory advertising and business practices.” The second open question concerns how ad targeting will interact with the European GDPR rules on consent; the third concerns how fair lending laws will consider ad targeting on variables that are highly correlated with protected classes but not the protected classes themselves.

“Ad delivery” is the term that explains the mechanical auction process by which algorithms allocate advertising space. The fact that “female eyeballs are more expensive” (women both control a greater share of household expenditures and have higher viewing-to-purchase rates [Cowgill and Tucker, 2020; Lambrecht and Tucker, 2019]) implies that ads from other retailers,

¹⁹ See Ali et al. (2019) for a detailed description of Facebook’s advertising platform.

²⁰ The relevant laws are the Fair Housing Act, the Equal Credit Opportunity Act (ECOA), Title VII of the Civil Rights Act, and the Age Discrimination in Employment Act. A creditor is allowed to engage in affirmative advertising to solicit or encourage traditionally disadvantaged groups to apply for credit (staff commentary to Regulation B, which implements ECOA).

²¹ In March 2019, Facebook settled five lawsuits that alleged that its advertising platform allowed companies to illegally target ads (Murphy, 2019), and in July 2019 the Equal Employment Opportunity Commission found that several employers violated federal law when they targeted their job ads on Facebook only to younger or male individuals.

such as employment opportunities in STEM fields in the Lambrecht and Tucker paper and presumably financial services as well, will be crowded out (Lambrecht and Tucker, 2019).²² Ali, Sapiezynski, Bogen, Korolova, Mislove, and Rieke (2019) conducted experiments that documented that Facebook was more likely to direct ads that included stereotypical male and female pictures to male and female audiences, respectively, even though the authors did not target their ads based on gender.²³ As Lambrecht and Tucker note, one way to counteract these distortionary effects of ad delivery is to target ads deliberately to different groups. This remedy, however, has the potential itself to be judged illegal discrimination for certain types of ads.

Discrimination Gateway #5: Practices of technology-facilitated advertising, including ad targeting and ad delivery, can lead to discrimination.

5 Innovation and Inclusion

The prior sections considered the role of technology in mitigating or facilitating discrimination resulting from human and algorithmic decisions. In this final, brief section, we step back from decisions and consider a final point that innovation itself is not immune from discrimination frontiers.

Worldwide, disparities exist in access to and inclusiveness of financial services. For example, in the U.S., younger, Black, and Latinx individuals are less likely to have a banking relationship (FDIC, 2020), and more likely to have a limited or insufficient set of credit history information such that access to credit is inhibited (Brevoort, Grimm, and Kambara, 2015); Black and Latinx individuals are more likely to have their mortgage applications rejected (Dietrich, Liu, Skhirtladze, Davies, Jo, and Candilis, 2019). Technology can reduce disparities in financial access and inclusion, such as through new modes of financial service delivery.²⁴ It can also increase these disparities,

²² Ali et al. (2019) subsequently verified the Lambrecht and Tucker finding by running a series of advertisements on Facebook and varying the maximum bid amount. The share of women shown the ad rose with the daily ad budget.

²³ Employment ads for lumberjacks were delivered primarily to white men, whereas black women were more likely to see employment ads for janitor positions.

²⁴ The widespread adoption of smartphones has increased access to financial services for individuals who live in communities underserved by traditional financial institutions. In the U.S., where 80% of adults own a mobile phone, the share of individuals owning a mobile phone does not vary by race (Pew Research Center, 2019), implying that innovations in mobile financial services have the potential to reach all racial groups equally. In 2016, U.S. regulators acknowledged the importance of access to technology for financial inclusion when they finalized rules that included “improving broadband access” as among the activities that might count toward a bank’s assessment under the Community Reinvestment Act. (See https://www.ffiec.gov/cra/pdf/2016_QA_Federal_Register_Notice.pdf).

such as when more privileged households are more likely to benefit from new technology.²⁵ The common perception (although untested and unclear) is that on the net, technology reduces disparities by expanding access.

Yet the impact on disparities is not equivalent to the impact on discrimination. Instead, each new innovation hits the market with the possibility that *even if the innovation reduces disparities*, it must be provided without any of the forms of discrimination discussed in this chapter. This can be a tall order, given the complex issues discussed in this chapter, and reaching those previously excluded from markets can be costly.

This creates a bit of an ironic outcome. Ex ante, discrimination may have caused disparities in access to financial services for some groups. However, innovators cannot simply use technology tools that remedy disparities without considering whether those tools discriminate against some groups. Regulators are keenly aware of this issue, and are increasingly providing support to help firms understand the fair lending implications of their innovations. For example, the Financial Conduct Authority (FCA) in the United Kingdom and the Consumer Financial Protection Bureau (CFPB) and the Office of the Comptroller of the Currency in the U.S. have introduced offices of innovation. The FCA and the CFPB have created “sandboxes” that allow companies to obtain regulatory relief for a limited period of time while the companies test financial innovations. The companies typically supply data in exchange and are monitored closely during the test period. In the United Kingdom, initial data suggests that the FCA’s Innovate program has reduced the time it takes for innovations to reach the market (Financial Conduct Authority, 2019).

6 Conclusion

The punchline of our analysis is that technological innovation has changed the way that discrimination manifests itself in financial services markets, with an ambiguous effect on the overall level of discrimination going forward.

²⁵ In the U.S., borrowers from peer-to-peer lenders appear to be more affluent than borrowers overall (Morse, 2015; DiMaggio and Yao, 2019). Likewise, mortgage borrowers who obtain their loans from FinTech firms tend to be more educated than mortgage borrowers overall (Fuster, Plosser, Schnabl, and Vickery, 2019). Dettling, Goodman, and Smith (2018) find that the spread of broadband technology led to students performing better on the SAT and applying to more colleges, presumably because broadband reduced the effort and informational costs of college applications. However, the effects were strongest for students from families of higher socioeconomic status, “suggesting that the technology may have increased pre-existing inequalities.”

The fact that technology has changed the manifestation of discrimination may not be a surprise. Yet once one starts to amalgamate all the channels of removing discretionary discrimination and increasing access to financial services while potentially introducing new discrimination through algorithmic process and human involvement in data and coding, the landscape shift becomes profound. Taste-based discrimination surely is less of a factor when decision-making in financial services markets is guided by algorithms rather than humans. However, opaque and complex algorithms drawing upon thousands of variables are increasingly governing how the creditworthiness of consumers is assessed, which advertisements consumers see, and which prices they are quoted. These processes have the potential to engender significant illegal statistical discrimination.

Regulation is evolving with technology, leading to the regulatory frontlines that we draw out in this chapter. Policymakers, regulators, and the courts must figure out how to modify the existing fair lending (and other equality law) infrastructure to a world where bias is embedded in data sets and the inner workings of algorithms. They must consider whether the traditional input-based framework for detecting and policing discrimination can be adapted to a world where the inputs have become so complex. They must judge when a variable is correlated enough with a protected class to be discriminatory itself, a task made more complicated by the sheer number of variables under consideration in algorithmic decision-making. And these decision-makers must also discern how the existing laws interact with new laws, such as the GDPR and other data privacy laws, as well as with public accommodation laws that have not historically been used to police discrimination in online spaces. The implications of these laws for technological innovation by financial services firms is undetermined as of the writing of this chapter.

Ideally, these momentous decisions would be aided and influenced by economic and policy research. We are concerned, though, that an inadequate amount of research will be produced that focuses on the welfare perspective of consumers of financial services. In our preparation of this manuscript, the imbalance away from such studies was evident. Because the algorithms are complex and the datasets are proprietary, much research, by necessity, is conducted in partnership with technology companies. This research is valuable, but the imbedded incentive structure will likely result in research that focuses more on the questions of interest to companies than to consumers.

We end finally on a positive note. The technological transformation of financial services is making enormous inroads into improving choices, competition, and access for millions of people. And although we indeed

conclude with the surprising thought that this technological transformation may do the opposite of leveling the playing field, this outcome is not set in stone. The choices made in the next few years will influence whether discrimination is a pervasive feature in the markets for financial services or not, and whether technological innovation results in these markets becoming more fragmented or more inclusive. The goal of our chapter is to bring these fault lines forward to improve the chances of reaching the inclusive outcome.

Acknowledgements The views in this chapter are not necessarily those of the U.S. Treasury Department, the Federal Reserve Board or its staff. Any errors are the sole responsibility of the authors. We are grateful to Katrina Blodgett, Carol Evans, and Varda Hussain for their careful review of the chapter and thank Bobby Bartlett, Tobias Berg, David Cross, Tim Lambert, Miles Larbey, David Palmer, Raghavendra Rau, Bradley Schnarr, Christopher Shelton, Luigi Zingales, and seminar participants at the University of Virginia for helpful comments and conversations.

References

- Alesina, Alberto F., Francesca Lotti, and Paolo Emilio Mistrulli. 2013. Do women pay more for credit? Evidence from Italy. *Journal of the European Economic Association* 11 (1): 45–66.
- Alexandrov, Alexei and Sergei Koulayev. 2018. No shopping in the U.S. mortgage market: Direct and strategic effects of providing information. Consumer Financial Protection Bureau Office of Research Working Paper Series, 2017-01.
- Ali, Muhammad, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through Optimization: How facebook's ad delivery can lead to skewed outcomes. [arXiv:1904.02095](https://arxiv.org/abs/1904.02095) [cs.CY].
- Altenburger, Kristen M. and Daniel E. Ho. 2018. When algorithms import private bias into public enforcement: The promise and limitations of statistical debiasing solutions. *Journal of Institutional and Theoretical Economics* 174 (1).
- Arrow, Kenneth J., 1973. The Theory of Discrimination. In *Discrimination in Labor Markets*, eds., Orley Ashenfelter and Albert Rees. Princeton, NJ: Princeton University Press.
- Athey, Susan and Michael Luca. 2019. Economists (and Economics) in Tech Companies. *Journal of Economic Perspectives* 33 (1): 209–230.
- Avery, Robert B., Mary F. Bilinski, Audrey Clement, Tim Critchfield, Samuel Frumkin, Ian H. Keith, Ismail E. Mohamed, Forrest W. Pafenberg, Saty Patrabansh, and Jay D. Schultz. 2018. A profile of 2016 mortgage borrowers: Statistics from the national survey of mortgage originations. National Mortgage Database Technical Report 6.0, April 18.

- Ayres, Ian. 1991. Fair driving: Gender and race discrimination in retail car negotiations. *Harvard Law Review*, 104 (4): 817–872.
- Ayres, Ian, Gary Klein, and Jeffrey West. 2017. The rise and (potential) fall of disparate impact lending litigation. In *Evidence and Innovation in Housing Law and Policy*, eds., Lee Anne Fennell and Benjamin J. Keys. Cambridge: Cambridge University Press.
- Barocas, Solon, and Andrew D. Selbst. 2016. Big data's disparate impact. *104 California Law Review*, 104: 671–732.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2019a. Consumer Lending Discrimination in the FinTech Era. Working paper.
- Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2020b. Algorithmic accountability: A legal and economic framework. Working paper.
- Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross. 2017. What drives racial and ethnic differences in high-cost mortgages? The role of high-risk lenders. *Review of Financial Studies*, 31 (1): 175–205.
- Becker, Gary. 1957. *The economics of discrimination*. Chicago: University of Chicago Press.
- Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri. 2020. “On the Rise of FinTechs: Credit Scoring using Digital Footprints.” *Review of Financial Studies* 33(7): 2845–2897.
- Berg, Tobias, Valentin Burg, Ana Gombovic, and Manju Puri. forthcoming. On the rise of FinTechs - credit scoring using digital footprints. *Review of Financial Studies*.
- Bhutta, Neil, Andreas Fuster, and Aurel Hizmo. 2019. Paying too much? Price dispersion in the US mortgage market. Working paper.
- Björkegren, Daniel, and Darrell Grissen. 2019. Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment. *World Bank Economic Review*. <https://doi.org/10.1093/wber/lhz006>.
- Blodgett, Su Lin, and Brendan O'Connor. 2017. Racial disparity in natural language processing: A case study of social media African-American english. [arXiv:1707.00061v1](https://arxiv.org/abs/1707.00061v1) [cs.CY] 30 Jun 2017.
- Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope. 2019. Inaccurate Statistical Discrimination. Working paper.
- Buolomwini, Joy, and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of Machine Learning Research* 81: 1–15.
- Bevoort, Kenneth, Philipp Grimm, and Michelle Kambara. 2015. Data Point: Credit Invisibles. Consumer Financial Protection Bureau Office of Research, May.
- Caliskan, Aylin, Joanna Bryson, and Arvind Narayanan. 2017. Semantics Derived Automatically from Language Corpora Contain Human-like Biases. *Science* 356 (6334): 183–186.
- Charles, Kerwin Kofi, and Erik Hurst. 2002. The transition to home ownership and the black-white wealth gap. *Review of Economics and Statistics* 84 (2): 281–297.

- Cowgill, Bo, and Tucker Catherine E. February 14, 2020. "Algorithmic Fairness and Economics." *Columbia Business School Research Paper*, Available at SSRN: <https://ssrn.com/abstract=3361280> or <https://doi.org/10.2139/ssrn.3361280>.
- Dastin, Jeffrey. 2018. Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*, October 9.
- Deku, Solomon Y., Alper Kara, and Philip Molyneux. 2016. Access to consumer credit in the UK. *European Journal of Finance* 22 (10): 941–964.
- Dettling, Lisa J., Sarena Goodman, and Jonathan Smith. (2018). Every little bit counts: The impact of high-speed internet on the transition to college. *The Review of Economics and Statistics* 100 (2): 260–273.
- Dietrich, Jason, Feng Liu, Akaki Skhirtladze, Misha Davies, Young Jo, and Corinne Candilis. 2019. Data point: 2018 mortgage market activity and trends. Consumer Financial Protection Bureau, October.
- DiMaggio, Marco and Vincent Yao. 2019. Fintech borrowers: Lax-screening or cream-skimming? Working paper.
- Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. forthcoming. "Measuring Bias in Consumer Lending," *Review of Economic Studies*.
- Doleac, Jennifer L., and Luke C.D. Stein. 2013. The visible hand: Race and online market outcomes. *Economic Journal* 123 (572): F469–F492.
- Donnelly, Robert, Francisco R. Ruiz, David Blei, and Susan Athey. 2019. Counterfactual inference for consumer choice across many product categories. [arXiv: 1906.02635](https://arxiv.org/abs/1906.02635) [cs.LG].
- Edelman, Benjamin, and Michael Luca. 2014. Digital discrimination: The case of Airbnb.com. Harvard Business School Working Paper 14–054.
- Edelman, Benjamin, Michael Luca, and Dan Svirsky. 2017. Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics* 9 (2):1–22.
- Edelman, Benjamin, and Abby Stemler. 2018. From the digital to the physical: Federal limitations on regulating online marketplaces. *Harvard Journal on Legislation* 56 (1): 141–198.
- Evans, Carol. 2017. Keeping Fintech fair: Thinking about fair lending and UDAP risks. *Consumer Compliance Outlook* 2: 1–9.
- Evans, Carol and Westra Miller. 2019. From catalogs to clicks: The fair lending implications of targeted, internet marketing. *Consumer Compliance Outlook* 3: 1–9.
- Federal Deposit Insurance Corporation (FDIC). October 2020. How America Banks: Household Use of Banking and Financial Services, 2019 FDIC Survey.
- Federal Trade Commission. 2016. Big data: A tool for inclusion or exclusion? Understanding the issues. Federal Trade Commission Report.
- Ficklin, Patrice and Paul Watkins. 2019. An update on credit access and the bureau's first no-action letter. Consumer Finance Protection Bureau Blog, August 6.
- Financial Conduct Authority. 2019. The impact and effectiveness of innovate. April.

- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. forthcoming. Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance*.
- Fuster, Andreas, Matthew Plosser, Phillipp Schnabl, and James Vickery. 2019. The role of technology in mortgage lending. *Review of Financial Studies* 32 (5): 1854–1899.
- Goodman, Bryce, and Seth Flaxman. 2017. European union regulations on algorithmic decision-making and a ‘right to explanation.’ *AI Magazine* 38 (3): 50–57. [arXiv:1606.08813](https://arxiv.org/abs/1606.08813) [stat.ML].
- Grundl, Serafin, and You Suk Kim. 2019. Consumer mistakes and advertising: The case of mortgage refinancing. *Quantitative Marketing and Economics* 17: 161–213.
- Gurun, Umit, Gregor Matvos, and Amit Seru. 2016. Advertising expensive mortgages. *Journal of Finance* 71 (5): 2371–2416.
- Hale, Brenda. 2018. Equality and human rights. Oxford Equality Lecture 2018. October 29.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. Human decisions and machine predictions. *Quarterly Journal of Economics* 133 (1): 237–293.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Cass R. Sunstein. 2018. Discrimination in the age of algorithms. *Journal of Legal Analysis* 10: 113–174.
- Lambrecht, Anja, and Catherine Tucker. 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science* 65 (7): 2966–2981.
- Morse, Adair. 2015. Peer-to-peer crowdfunding: Information and the potential for disruption in consumer lending? *Annual Review of Financial Economics* 7: 463–482.
- New York Department of Financial Services. 2019. Governor Cuomo calls on DFS to investigate claims that advertisers use facebook platform to engage in discrimination. Press release, July 1.
- Perry, Vanessa G., and Carol M. Motley. 2009. Reading the fine print: Advertising and the subprime mortgage crisis. *California Management Review* 52 (1): 1–16.
- Pew Research Center. 2019. Mobile technology and home broadband 2019. June 13.
- Phelps, Edmund. 1972. The statistical theory of racism and sexism. *The American Economic Review* 62 (4): 659–661.
- Pop, Valentina. 2013. German data chief attacks credit-profile firms. EU Observer, April 25. <https://euobserver.com/economic/119930>.
- Pope, Devin G., and Justin R. Sydnor. 2011. Implementing Anti-Discrimination Policies in Statistical Profiling Models. *American Economic Journal: Economic Policy* 3: 206–231.
- Ravina, Enrichetta. 2019. Love and loans: The effect of beauty and personal characteristics in credit markets. Working paper.

- Roth, Lorna. 2009. "Looking at Shirley, the Ultimate Norm: Colour Balance, Image Technologies, and Cognitive Equity," *Canadian Journal of Communication* 34: 111–136.
- Rothstein, Richard. 2017. *The color of law*. New York: Liveright Publishing Corporation.
- Safronova, Valeriya. 2019. Gender stereotypes banned in British advertising. *New York Times*, June 14. <https://www.nytimes.com/2019/06/14/style/uk-gender-stereotype-ads-ban.html>.
- Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Russo. 2003. Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities? *Quantitative Marketing and Economics* 1 (1): 65–92.
- Stowe, Stacey. 2019. New York City to ban discrimination based on hair. *New York Times*, February 18. <https://www.nytimes.com/2019/02/18/style/hair-discrimination-new-york-city.html>.
- Turner, Lee N. 2018. "Detecting racial bias in algorithms and machine learning", *Journal of Information, Communication and Ethics in Society* 16 (3): 252–260. <https://doi.org/10.1108/JICES-06-2018-0056>.
- Woodward, Susan E. and Robert E. Hall. 2012. Diagnosing consumer confusion and sub-optimal shopping effort: Theory and mortgage-market evidence. *American Economic Review* 102 (7): 3249–3276.



Trust Building and Credit Reporting with Big Data in the Digital Age

Wensheng Peng and Feng Zhu

1 Introduction

Trust is the foundation of human society and essentially all social interactions, especially commercial and financial activities. Confucius said over two thousand years ago: “Without trust, a society cannot stand” (民无信不立). Many transactions are made under imperfect information, with either the buyer, the seller, or both, being uncertain about the quality of the traded goods or traders themselves. The need for trust is greatest where information asymmetry is most severe, namely insurance and credit intermediation. Over the history, many different trust building mechanisms have been designed, and credit reporting emerged in the nineteenth century has become an indispensable component of modern financial infrastructure. An effective credit reporting and scoring system is the cornerstone of commerce and finance, it helps smooth commercial transactions and financial operations, thereby supporting economic growth. Today, with new technologies,

F. Zhu (✉)

Senior Research Director, Luohan Academy, Hangzhou, China
e-mail: zhu.feng@antfin.com

W. Peng

Chief Economist, China International Capital Corporation Limited, Beijing, China
e-mail: wensheng.peng@cicc.com.cn

the digital economy has brought immense new opportunities as well as new challenges to trust building, and China is among the economies that have experienced the most sweeping changes.

Trust-building and credit reporting are fundamental to financial activities. As Stiglitz (2000) points out, “it is now recognized that information is imperfect, obtaining information can be costly, there are important asymmetries of information”. Traditional lenders often rely on various means including relationship banking or lending against collaterals to reduce information asymmetry and lower default risks and monitoring and screening costs. Unable to establish a long-term relationship with banks or to provide collaterals, many small businesses and households are denied access to credit. Credit reporting helps reduce information asymmetry and promote more efficient resource allocation, and a more developed credit scoring system enhances financial inclusion, by helping finance individuals and micro- and small and medium-sized enterprises (MSME) which are often in greatest need of funds. This is especially relevant in China, an economy with a sizeable financing gap. World Bank (2019b) reports that only 26% of the Chinese population over the age of 15 borrowed from financial institutions or used a credit card in 2017, well below the average of 45% among the Organization for Economic Co-operation and Development (OECD) countries.

An efficient credit scoring system is part of a broad trust building system that promotes trust and supports growth by reducing the operating costs of the entire social and economic system. For many years, inadequate incentives and impunities for breaches of trust and dishonest behavior in China have added substantial costs to economic activities. This hurts commerce and finance, making it difficult to attract foreign as well as domestic investment. The rapid development of internet, mobile communications, and especially the rise of new technology companies and data analytics, have brought new opportunities but also challenges to trust building and credit reporting. In the digital age, the emerging platform economy brings together billions of strangers to trade with each other on a wide range of platforms including Alipay, Amazon, Facebook, ebay, JD.com, LendingClub, Mercado Libre, Taobao, Tmall, and WeChat. For instance, two parties of a Taobao transaction are, on average, over 1,000 kilometres apart. New mechanisms have been designed to provide trust among strangers to enable transactions, early applications included reputation-building mechanisms such as feedbacks from buyers and sellers and escrow services such as Paypal and Alibaba Secure Payment.

Growing trades and social interactions have provided unprecedented access to data. Opportunities abound. First, traditional credit reporting mostly relies

on historical loan data of financial institutions, and information exchange among different credit scoring entities is often inadequate. The digital age has brought with it a greatly expanded range of data sources covering a diverse set of information. While traditional models typically focus on credit history, especially bank and credit card lending, new models have significantly expanded the application scenarios and ecosystems well beyond financial services, retrieving information from “non-traditional” individual data on employment, taxation, online and offline commercial and financial transactions, social networking, and a myriad of other activities. Second, technological advances have dramatically lowered the costs of data collection, processing and storage. Big data analysis based on cloud computing has led to a dramatic rise in data-based decision making, and the enhanced ability and predictive accuracy of trust building and credit reporting systems allow real-time assessment of the reputation and credit standing of individuals and firms. The use of automated algorithms expedites the scoring process. Third, the credit reporting industry has since undergone a major change globally, and in China in particular.

The chapter is organized as follows. Section 2 provides an overview of major credit scoring models, emphasizing the experiences of China, the United States, Europe, and Japan. Section 3 discusses changes and challenges brought to trust building in the new digital age, against the background of major innovations in technology and business models. China’s recent experience is especially relevant, and a notable case has been the emergence of Zhima (Sesame) Credit and other similar agencies. Section 4 concludes.

2 Traditional Trust Building

A practical approach to trust building is credit reporting focussed on facilitating financial transactions. Durand (1941) was among the first to suggest the use of statistical methods, and credit rating formulae, to discriminate between good and bad loans with credit scores.¹ The rising use of credit cards since the late 1960s has led banks and other credit card issuers to look for solutions to cope with the growing needs for automated lending decision, and credit scoring gained importance, especially after the implementation of the Equal Credit Opportunity Act enacted in the United States in October

¹ Curiously, one of the earliest uses of such scoring was by Sears in the 1950s to decide to whom it would send its catalogues. Similar uses of scoring for non-credit-related areas seem to have picked up in China, e.g. the use of Ant Group’s Zhima Scores to provide deposit-free bike or power bank rental services.

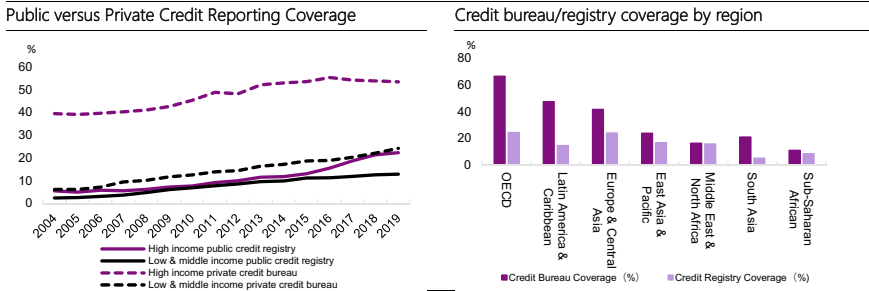


Fig. 1 Growing Credit Reporting Coverage (*Note* Coverage by public credit registry and private credit bureau in low- and middle-income economies versus the coverage in high-income economies. *Source* World Bank)

1974. In the 1980s, the use of credit scoring spread from credit card lending to other bank loans.² Credit scoring has since been used in almost all forms of lending and beyond: credit cards, personal loans, car finance, insurances, utility payments, mortgage lending, and small business loans. At the same time, behavioral scoring systems were built to allow lenders to better manage existing clients by forecasting their future performance, making informed decisions on credit limits, the clients targeted for marketing new products, and the recovery of the nonperforming debt.

Over the past fifteen years, public and private credit reporting coverage has risen steadily across the world (see Fig. 1). The coverage has stayed significantly higher in high-income economies, where there is a long and venerable tradition of credit reporting, in typically more mature credit markets and with better data. The average credit registry coverage of the adult population in low-income economies is less than 3%, compared to over 22% in high-income ones. Comparing different groupings, OECD countries enjoy the highest coverage by credit bureaus and registries, while Sub-Saharan Africa has the lowest coverage. In emerging economies, credit registries and bureaus also tend to collect less comprehensive information, this limits businesses' access to credit and hurts economic performances. World Bank (2019a) suggests "notable differences between stronger and weaker performing economies" in the quality of regulation and information: in the top-20 ranked economies, 83% of the adult population on average is covered by either a credit bureau or registry, whereas in the bottom 50, the coverage is only 10%.

² Olegario (2003) provides an overview of the historical evolution of credit reporting agencies, especially in the United States.

2.1 Major Credit Scoring Models

Global credit reporting systems can be classified into three different representations, namely the market-oriented US model, the public model most common in Europe, and the industry association model of Japan. The US credit reporting system consists of credit bureaus, i.e. profit-making credit scoring companies that independently collect and process personal and corporate credit data. The system is dominated by three major credit bureaus (Experian, Equifax, and TransUnion), and known for its time-tested personal information interaction mechanisms, rigorous data processing, and sound regulatory measures. Lenders such as banks and credit card companies use credit scores as a major reference in assessing potential loan risks. To evaluate personal creditworthiness, Fair Isaac Corporation (FICO) scores are used by most credit scoring companies. While consumer credit bureaus typically use FICO's algorithm to compute credit scores, for mainstream corporate credit scoring agencies such as Dun & Bradstreet, Equifax, and TransUnion, there are no universally accepted industry standards for calculating business credit scores.

In Europe, credit scoring agencies are broadly divided into public credit registries and private credit bureaus, and in some countries, the credit reporting system relies almost entirely on credit registries. Public and private credit agencies operate in a similar fashion, except for the mandatory nature of data collection in the former case. Public credit registries are administered by the respective central banks, to which only authorized central bank staff and reporting financial institutions have access. Financial institutions must proactively provide personal loan data to credit registries, which then consolidate data for individual borrowers from various banks before sending the consolidated loan data back to the banks. The focus is on the overall financial situation of a borrower, and his credit report contains both positive and negative information (e.g. default and arrears statements). In general, credit registries are characterized by mandatory participation, confidentiality and privacy protection, reporting thresholds for loan information, and computer-intensive technologies. In Germany, the threshold is so high that credit registries only serve large borrowers, while it is very low in France and Belgium, covering many home loan applicants.

Led by industry associations, credit reporting in Japan started in 1892, and it functions as market-oriented operations. Personal credit scoring is membership-based with three different layers, namely banking, consumer credit, and sales credit, each associated with the Japanese Bankers Association, the Japan Consumer Credit Industry Association, and the Federation

of Credit Bureaus. Each association has established its own credit information center. Initially, the information centers were relatively independent with segmented operations, interactions and information sharing were confined to own industry. In 1987, with the support of the three associations, the Credit Information Network System, a cross-industry network, was established to encourage information sharing. Japan's credit scoring has become dominated by Teikoku Databank and Tokyo Shoko Research. Teikoku Databank commands over 70% of Japan's business credit scoring market, it is one of Asia's largest business credit database, with information on 4,000 listed companies and 2.3 million non-listed companies.

2.2 Credit Reporting with Chinese Characteristics³

Credit reporting in China started with the creation of Far East Credit Rating Co., Ltd (远东资信评估有限公司) in February 1988, in Shanghai. The country's credit reporting system has its origin in the "Loan Certification System" from the early 1990s, which led to the construction of the "Bank Loan Registration and Inquiry System" in 1997–2002. In 1999, China launched a pilot personal credit scoring program, managed by Shanghai Credit Information Services Co., Ltd. This was followed by other local initiatives in Shenzhen, Dalian, Beijing, and other cities.

Credit reporting had a major boost in 2004–2006, with the issuance of new guidelines on personal and enterprise credit reporting. Under the tutelage of the People's Bank of China (PBOC), banks set up a nationwide, centralized enterprise and personal credit reporting system that became operated and managed by the PBOC Credit Reference Center (CRC) established in 2006 (see Fig. 2).⁴ The CRC has grown steadily over time, it is now considered the world's largest credit reporting system in terms of coverage. As of May 2019, the CRC collected information on 990 million individuals and 25.9 million firms and other institutions, including 13.7 million small and micro-businesses. On average, the CRC provided assistance to about 5.5 million and 300,000 inquiries per day on the credit standing of individuals and firms, respectively.

³ Much has been written on China's credit reporting system. Huang, Lei and Shen (2016), and Chorzempa (2018) provide a recent overview of its evolution. More recent developments in China's social credit system have attracted attention, e.g. Backer (2019), Chorzempa, Triolo and Sacks (2018), Liu (2019), Ohlberg, Ahmed and Lang (2017), Sithigh and Siems (2019), and Zhang, He, Liu, Chen, and Liu (2019). The system has a broad coverage that includes political, social, judicial, and commercial trusts, and is still under early construction.

⁴ The System is officially known as the first-generation "Financial Credit Information Foundational Database" (金融信用信息基础数据库).

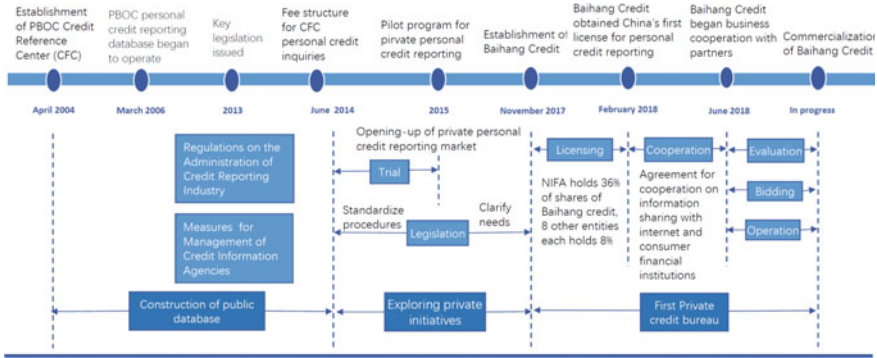


Fig. 2 Recent development of the credit reporting system in China

The landmark “Regulations on the Administration of Credit Reporting Industry” released in early 2013 stipulated that the system is a state-administered credit information database, providing the legal and regulatory framework for the industry (see Fig. 3). Under the Regulations, three types of institutions are requested to provide information: commercial and policy banks; non-bank financial institutions including trust companies, finance companies, and auto finance companies; and other institutions including micro-credit companies, financing guarantee companies, and housing provident fund management centers. The database collects corporate and personal information covering finance, social security, provident funds, tax arrears, court rulings, and telecommunication and public utilities payment records.

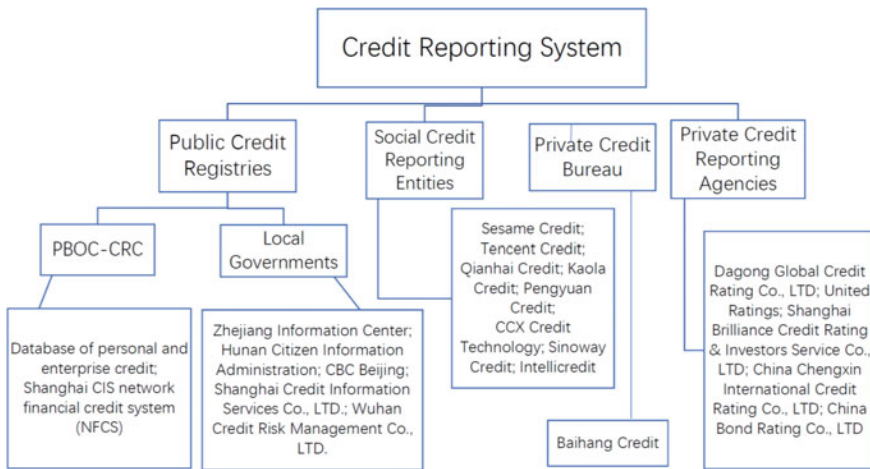


Fig. 3 Current personal credit reporting system in China

The information has proved to be very useful. In 2018, 21 banks relied on the system to turn down loan applications worth CNY 911.7 billion, issued advance warnings on outstanding loans worth CNY 1.3 trillion, and cleared up CNY 159.4 billion in nonperforming loans. By the end of 2019, the personal and corporate system each had access to 3,737 and 3,613 institutions, respectively, basically covering all types of formal lenders, and the “first-generation” system had collected data on 1.02 billion persons, 28.341 million enterprises and other organizations.

On January 19, 2020, the PBOC Credit Reference Center launched its second-generation credit reporting services. Compared with its previous version, the new system has improved information collection, product design, technical infrastructure, and data security. The information content has been enriched to better reflect the subjects’ credit profile. The CRC now collects information on revolving loans, large credit card installments, joint borrowers, corporate guarantees for individuals and personal guarantees for firms, and payment of debt overdue.

The current credit reporting system combines government-led initiatives with private, third-party entities including Baihang Credit and other credit reporting companies (Fig. 3). The 2013 Regulations allow the establishment of commercial credit reporting companies, which should register with an appropriate PBOC branch. As of November 2019, a total of 135 commercial credit bureaus in 22 provinces and municipalities had completed their filings with PBOC branches, a significant number of these located in Beijing (38) and Shanghai (29), but only 10 in Guangdong, then the largest provincial economy in the country.

Despite significant progress in the last decade, China’s credit reporting system still lags far behind the more mature economies. Several factors may explain this. First, the demand for credit reporting only emerged with the growth of more market-oriented business entities following China’s economic reforms launched in 1978, the operating, legal and regulatory framework for credit reporting took time to evolve. Second, market-oriented personal credit reporting facilities in China are still in an early stage of development, so far only Baihang Credit obtained a license, with a limited product line. Despite a large number of existing commercial credit reporting companies, they are generally small in size and the scope of services is limited and inadequate considering the needs of a fast growing economy. The existing legal and regulatory framework also needs to be strengthened to enhance privacy and consumer protection, a key factor in promoting the healthy growth of the industry.

In addition, the CRC mainly collects financial data nationwide, its coverage of the non-financial data is limited, and it lacks a more timely update of information. Most other credit agencies work in separation, with little or no information exchange and sharing, and inadequate and segmented information makes it hard for the industry to progress and achieve scale efficiency. The emergence of big data analytics, along with the rise of technology companies and their new business models, presents new opportunities as well as challenges in breeding better coordination mechanisms.

3 Trust Building in the Digital Age

The advent of the digital age, made possible by a myriad of technological advances, has fundamentally altered the way how people interact with each other and conduct businesses. Internet and mobile communication technologies have provided the infrastructure enabling the rapid growth of multi-sided platforms on which social and economic interactions take place, typically among complete strangers or distant traders. While the problem of the lack of trust inhibiting trades is not new, it turns out to be much more debilitating for online platforms, requiring novel and better solutions. The growing flows of remote online transactions, financial or commercial, depend on mutual trust, and many of the successful platforms have also been vanguards in devising mechanisms that build and maintain trust underlying such flows. These platforms include social networking and messaging apps, e.g. DingTalk, Facebook, Instagram, Twitter, WeChat, WhatsApp, as well as e-commerce giants Alibaba (e.g. Taobao and Tmall), Amazon, eBay, JD.com and Mercado Libre.

Trust is crucial, without which no deals can be completed. To cope with the mounting challenges from growing online transactions, many platforms have engineered multifold and multifaceted solutions. Trust building is more encompassing and typically relies on more varied information sources than credit reporting, but both work to reduce incentives for devious behavior often originating from information asymmetry. Stiglitz (2000) emphasizes the central role of mechanisms like reputation in a world of imperfect market for information. Online marketplaces like Alibaba, Amazon, and eBay all have reputation mechanism such as rating systems encouraging buyers and sellers to provide information on each other. Resnick and Zeckhauser (2002) use 1999 eBay data to study its which has a feedback system collecting comments from buyers and sellers about each other after each transaction. They find that reputation profiles were predictive of future performance, and the

players reciprocate and retaliate. Luohan Academy (2021) describes Taobao's "Detailed Seller Ratings" system designed to improve the exchange of information and promote trust, which played an important role in promoting firm growth.

In fact, trust deficit motivated the launch of Alipay's escrow service in 2004, enabling the very first trade on Alibaba's Taobao marketplace. On October 18, 2003, Zhengzhong Jiao, a student in Xi'an, China, successfully bid on Taobao a second-hand Fuji camera worth CNY 700 from Weiping Cui, a student in Yokohama, Japan. Online C2C transactions then took place as trades in close proximity or remote trades via bank transfers of funds. But Mr Jiao was unwilling to transfer money without receiving and checking the quality of the camera, while Mr Cui would not mail it without being paid first. The deal almost collapsed due to a lack of mutual trust. By chance, Mr Cui came across Taobao's "secured transaction" service, through which a seller initiates a transaction, the buyer pays, then the service provider places the fund in safekeeping. Once the buyer confirms the receipt of the merchandise, his payment is forwarded to the seller. Mr Cui became the very first user of this service, which later gave rise to Alipay, now a major player in third-party retail payment in China.

Credit reporting faces major challenges in the digital age, with a massive increase in decentralized, remote and anonymous interactions.⁵ This affords credit reporting agencies possibilities of novel and more powerful solutions.

3.1 Technological Advances, Data Analytics, and Credit Scoring

Information has always been a "fundamental component of all financial transactions and markets" (Liberti and Petersen 2019), and indeed of all commercial transactions. The value of information is greater where the problem of information asymmetry is most severe, notably in lending and insurance businesses, or in other typical markets for lemons (Akerlof 1970). Over the last two decades, technological advances and private arrangements of information sharing have fostered trust building and increased the use of credit scoring in almost all forms of loan origination. The growing amount of data and the ability to automate credit decisions have transformed previously local and personal markets into national or global, and increasingly impersonal markets. The wide use of electronic interfaces for transactions

⁵ Hurley and Adebayo (2017) provide a critical review and detailed discussion of credit scoring in the digital age, in particular new opportunities and possible regulatory challenges.

among strangers implies an increasing dependence of such transactions on trust building based on hard, quantifiable information.

3.1.1 Growing Use of Alternative, Big Data

For much of the human history, most information has been collected and decision made locally among people known to each other, where “soft information” played a prominent role.⁶ As information and communication technologies advanced, transactions among unknown market participants have become commonplace, the scope for quantifiable and verifiable “hard information” has grown rapidly. The quantitative nature of hard information suggests that it can be more easily collected, stored, and transmitted electronically, and soft information has also become easier to be transformed into numbers on a large scale. The accumulation and sharing of large amount of valuable quantitative or quantified information, allied with new business models, automated algorithms and decision-making have had a significant impact on credit reporting in China and elsewhere.

Traditional credit reporting typically focuses on five dimensions: individual character; capital; collateral or own resources; capacity or repaying ability; and market conditions. For example, FICO collects mainly financial information and focuses on factors in five areas to determine creditworthiness, namely, payment history (35%), debt burden (30%), length of credit history (15%), types of credit used (10%), and recently opened credit accounts (10%). FICO scores range from 300 to 850.

On the other hand, many newcomers take an “all data is credit data” approach. ZestFinance, a US firm founded in 2009, offers big data credit scoring tools to providers of payday loans to predict credit risk on the basis of non-traditional data. It combines traditional credit information with thousands of data points for each individual, collected from consumers’ offline and online activities, including the borrower’s data, its own proprietary data, public data, and social network data. Based on big data analytics, Alibaba’s Zhima trust scores offer external entities a means to evaluate social and commercial trustworthiness, a use far broader than mere credit scoring. Zhima scores each user on a scale of 350 to 950 points, based on a rich

⁶ Stein (2002) distinguishes information that is “soft”, or information not directly verifiable by anyone other than the agent who produces it; and “hard”, verifiable information, such as financial statements or tax returns. Liberti and Petersen (2019) provide detailed discussions on the evolution of these concepts. Berger, Miller, Petersen, Rajan, and Stein (2005) find evidence that small banks can better collect and act upon soft information than large banks which are less willing to lend to clients with less information, such as firms with no financial records.

set of weakly correlated data including his/her financial credit records along with payment history; personal information, and information on available assets, including the use of financial products and services and Alipay account balances; behavioral trends in commercial transactions; actual behavior and revealed preferences; and social relationships and interactions.

Barron and Staten (2003) show that lenders could significantly reduce default rates by including more comprehensive borrower information in their default prediction models. Nevertheless, the implicit “big data hubris” assumption that big data are a substitute for, rather than a supplement to, traditional data collection and analysis is wrong. As Lazer, Kennedy, King, and Vespignani (2014) point out, “implicit big data offer enormous possibilities for understanding human interactions at a societal scale, with rich spatial and temporal dynamics, and for detecting complex interactions and nonlinearities among variables”, yet “traditional ‘small data’ often offer information that is not contained (or containable) in big data, and the very factors that have enabled big data are enabling more traditional data collection”. In terms of credit reporting, alternative, or “hardened” soft data augments traditional data, and credit history and financial statements continue to be among the most valuable core information in assessing consumer default risks, and alternative data is most useful in cases where a customer has no or patchy credit history.

One important component of the alternative big data is the so-called “digital footprints”. These are trails users unintentionally leave online (passive), or intentionally submit (active). They include data that users create while visiting or registering on websites, blogging, tweeting, posting comments, sending emails, submitting information to online service providers, or simply viewing a page or clicking a specific link. For instance, ZestFinance takes the speed at which a loan applicant scrolls through an online terms and conditions disclosure as an indicator of how responsible she is. Focusing on a sample of 270,399 purchases from an e-commerce company in Germany, from October 2015 to December 2016, Berg, Burg, Gombović, and Puri (2019) study the informativeness of digital footprints for credit scoring. They find that “even simple, easily accessible variables from a digital footprint match the information content of credit bureau scores”. Digital footprints complement rather than substitute for core financial information, they affect credit access and reduce default rates.

The often profuse use in internal ratings by Chinese tech companies of big data originating from a wide range of alternative sources, including digital footprints, has become a common practice. This is particularly relevant for assessing the trustworthiness or creditworthiness of individuals and small

firms with little or no credit history. Based on transaction-level data from a Chinese fintech company in May–September 2017, Gambacorta, Huang, Qiu, and Wang (2019) test different models, they find that the model based on machine learning and non-traditional data performed better in predicting losses and defaults than traditional models in the presence of a negative shock to the aggregate credit supply. However, the comparative advantage of this model diminishes for borrowers with a longer credit history, suggesting that the marginal value of soft information declines with richer core or hard information.

The information sources of traditional models and new models based on big data analytics are rather different. Traditional credit reporting agencies collect mainly financial information, the coverage, and size of which are relatively limited. More importantly, data collection remains inadequate in China, and data sharing among various traditional credit reporting entities is far from being effective, this may lead to inaccurate or inconsistent ratings. Newer rating schemes, e.g. Ant Group's Zhima scores Sesame Credit, typically cover data of far greater dimensions. Yet, for the purpose of credit scoring, financial information remains most valuable; alternative data complements it and adds further value, especially for the evaluation of trustworthiness.

3.1.2 Big Data Analytics

Trust building and credit reporting assess the trustworthiness and creditworthiness based on past data, the ability and accuracy of which have been tremendously enhanced with the rise of data analytics. Data analytics are built around big data, using technologies involving cloud computing, artificial intelligence (AI), machine learning (ML), and automated algorithms that facilitate data analyses and decision-making. Big data, characterized by Laney's (2001) three v's, i.e. massive volume, variety of forms it takes, and velocity with which it arrives, has led to the creation of a new generation of data infrastructure supporting business and policy decisions.⁷ Typical examples include Equifax's NeuroDecision which uses Explainable Artificial Intelligence to generate "logical, actionable explanations" for consumer credit scoring; JPMorgan's Contract Intelligence (COIN) that automates document review for certain contracts; and Bloomberg's Social Sentiment Analytics (in partnership with Twitter).

⁷ Dietrich, Heller, and Yang (2015) provide a detailed overview of data analytics. Diebold (2019) digs into the origin and evolution of the term "big data".

What is new is the widespread applications, and the confluence or integration of the rapidly advancing digital technologies and financial and real activities, creating new business opportunities and new sources of data. As communication and storage costs decline, mobile and internet coverages rise, social, economic, and financial connectedness has grown at a fast pace (Fig. 4). The broadening and deepening socioeconomic interactions have led to a growing range of available data sources, facilitated by tech companies' new business models that emphasize ecosystems with diversified business scenarios where the economy of scope plays an important role. Some tech companies have gained ability to track and link digital behavior across platforms, websites, and online and offline activities for billions of users. Intensive networking and rich data allied with enhanced computing have generated significant growth in the use of hard information.

As Tirole (2020) points out, a credit rating system helps financial institutions to fend off bad borrowers, and big data analytics has enabled a more inclusive access to funding for Chinese SMEs. Chorzempa (2018) stresses the need to upgrade China's financial information systems in order to establish a vast number of missing credit histories, and make it easier to share the existing ones. The system should incorporate traditional credit data as well as alternative data such as telecom and rent payments to quickly absorb those who lack a traditional credit history.

For years, both Alibaba and Tencent's micro-credit services have relied on internal ratings based on big data analytics to make informed decisions on lending. Both firms are leading providers of data and technology services in

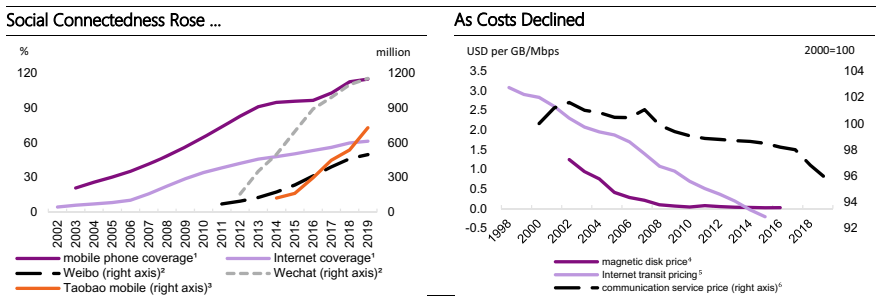


Fig. 4 Technological Advances Enhance Social Connectedness (Notes ¹Coverage ratio is calculated as the percentage of people using mobile phones or Internet. ²Weibo, Wechat are Chinese social apps. ³Taobao mobile is the mobile app of the largest e-commerce platform in China. ⁴Magnetic disk price is calculated as the unit price (USD per Gigabyte) of data storage. ⁵Internet transit prices are from drpeering.net and are US data. ⁶Communication service price is derived from China's CPI component. Sources Wind, hblok.net, drpeering.net)

China, with expanding business lines and increasing diversity of their various platforms that provide rich access to user data. Payment services such as Alipay or WeChat Pay serve both online and offline consumers and firms, including billions of active users on e-commerce giants JD.com, Taobao and Tmall; food deliveries Ele.me and Meituan Waimai; logistics firms Cainiao Smart Logistics Network and JD Logistic; social media Tencent QQ, WeChat and Weibo microblogging; and travel services such as Fliggy Travel, among others. Rich data afford tech companies further insights and opportunities to better serve their customers while expanding their ecosystems, which in their turn, render more and better data, forming a virtuous information loop. On the other hand, Alibaba, Baidu and Tencent's cloud services and their AI experts have provided the essential ability to structure and analyze the collected data to create value for varied uses including internal ratings for decisions on user access or on consumer loans (e.g. Huabei and Jiebei) and small business credit (e.g. MYBank and WeBank).

3.1.3 Artificial Intelligence Methods and Automated Algorithms

Credit reporting and its core methodology have been perfected over the years. Credit scoring applies statistical techniques to predict the probability of a loan applicant or an existing borrower defaulting or becoming delinquent. Typically, credit scoring is implemented using a scorecard that assigns points to key individual characteristics and transaction aspects to produce a numeric estimate of the individual's credit risks. Many analytic methods, such as regression analysis, classification, and cluster analysis, have been used in the past. Thomas (2000) reviews the main techniques, both statistical and operational research, used to support credit and behavioral scoring. Most scorecard builders use one of the techniques or a combination of them. The statistical tools include discriminant analysis which is essentially linear regression, or a variant called logistic regression, and classification trees (also known as recursive partitioning algorithms). The operational research methods include variants of linear programming.

Traditionally, discriminant analysis and logistic regression were among the most popular techniques. More recently, Artificial Intelligence (AI) methods more suitable for big data environments have been applied. Louzada, Ara, and Fernandes (2016) reviews 187 papers on credit scoring published in 1992–2015, which applied statistical techniques such as linear regression and logistic regression, discriminant analysis, as well as a number

of nonparametric statistical and AI modeling approaches, including artificial neural networks, support vector machine, decision trees, fuzzy logic, genetic algorithms and genetic programming, nearest neighbor methods, Bayesian networks, hybrid methods, and ensemble methods such as bagging, boosting, and stacking. Neural networks, support vector machine, hybrid and combined techniques are found to be most common, and support vector machine method appears to provide high predictive performance with low computational efforts.

Using US Survey of Consumer Finances data, Munkhdalai, Munkhdalai, Namsrai, Lee, and Ryu (2019) compare machine learning (ML) approaches, with regression models based on multivariate adaptive regression splines, support vector machine, random forest, extreme gradient boosting, and artificial neural network, with the traditional expert-based FICO credit scoring model. They find that if lending institutions in early 2000s had used credit scoring models based on ML methods, their expected credit losses would have been lower. In addition, deep neural networks and extreme gradient boosting algorithms achieve the best performance in terms of accuracy. AI/ML methods may improve credit scoring efficiency and support better lending decisions.

A growing number of tech firms have adopted the use of sometimes complex algorithms that help detect behavioral patterns and predict credit risks using big data analytics. This has proved to be beneficial especially for millions of individuals with “thin credit file” who otherwise would have no access to credit. Increasingly, automated algorithmic decision systems combine machine learning with traditional statistical methods to classify individuals’ credit risks to improve predictive accuracy and fairness. Einav, Jenkins, and Levin (2013) find the value of automated credit scoring for a large auto finance company which enabled changes in its lending practices. It appears to have raised profits by enhancing the ability to screen high-risk borrowers and the ability to target more generous loans to lower-risk borrowers.

3.1.4 Enhancing China’s Credit Reporting in the Digital Age

China’s market reforms and liberalization have been the driving force of rapid economic expansion. As the economy grows, China has seen more widespread and rising cases of economic crimes as well as degrading social behavior that defies traditional values and norms. As the economy expanded at rapid pace, credit demand also rose fast, and growth in all types of loans, especially MSME loans, has outpaced US loan growth since 2007 (see Fig. 5). Frauds

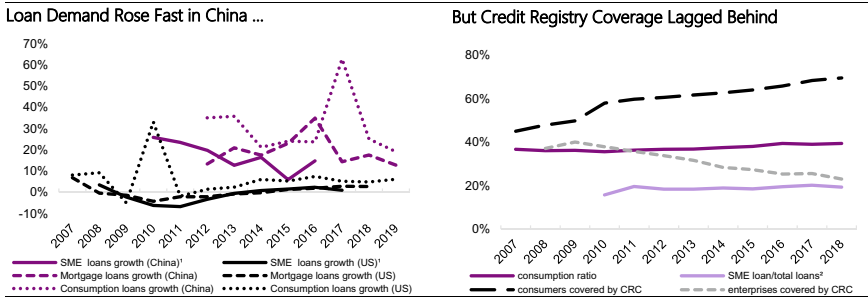


Fig. 5 Trust is Needed When Credit Demand Grows (Notes ¹Annual growth rates of SME, Mortgage, and Consumption loans. SMEs are classified according to OECD standards. ²Ratios of consumption to GDP and of SME loans to total loans. Statistics of SME loans come from PBOC and SMEs are classified according to the standards of the National Bureau of Statistics. Consumers covered by China's Credit Reference Center (CCRC) is calculated as the percent of total population. ³Scaled by the total population. ⁴Scaled by the number of market entities published by State Administration for Market Regulation. Sources OECD; PBOC; Wind)

and defaults proliferated. In the meantime, the credit registry coverage ratio for enterprises actually fell, highlighting a growing trust deficit. In response, policymakers released the 2014 plan to establish a comprehensive social credit system to promote trust among people, which motivated Tirole (2020). He suggests that social scores may enhance trust in society, and “they have already promoted better behavior on e-commerce and ride-hailing platforms around the world”.

As the digital economy grows, trust deficit has become more acute for online platforms. In the past decade, the ratio of online sales to total retail sales rose from almost zero in later 2000s to over 20% in the last few years, accompanied by a leap in third-party payments and a rise of loans by digital or virtual banks such as MYBank and WeBank (see Fig. 6). Notably, the trading volume of P2P lending platforms in China rose rapidly but began to fall following the failure of Ezubao in December 2015, and many more platforms failed after the 2016–2017 regulatory tightening, reflecting the acute problem of information asymmetry and trust deficit in these platforms due to a lack of relevant data and rampant fraudulent behavior.

Among the tech firms in China, Alibaba was the first to venture into trust building, and it has remained among those most actively engaged. As early as 2003, its e-commerce platform Taobao designed and implemented a simple three-level internal rating system based on customer feedbacks and other data such as delivery time and seller history. In 2004, the system was upgraded to a more elaborate system for all online stores, and gradually revamped to include detailed scores for services, logistics, and product quality into the “Detailed

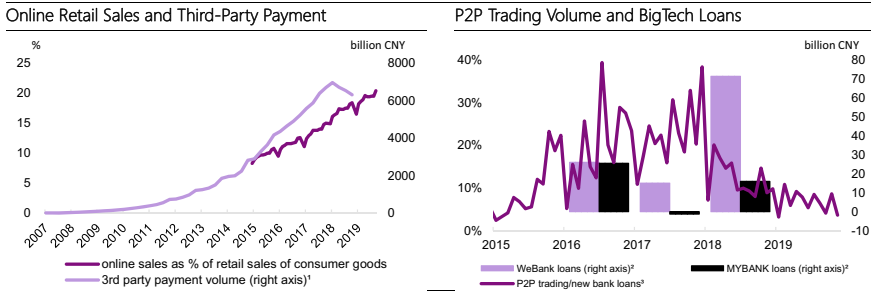


Fig. 6 E-commerce, Digital Payment and Digital Financing (*Notes* ¹Ratio of online sales as % of retail sales of consumer goods; total volume of third-party payment. ²WeBank and MYBANK new loans. ³Ratio of P2P trading volume to new bank loans. *Source* Analysys; PBOC; WDJ.com; Wind)

Seller Ratings” in 2009. The system expanded into a “Rating Social Network System” in 2012, incorporating further comments on user reviews and the Q&A section. The system has played an important role in promoting firm growth, as sales tend to rise significantly after a rating upgrade, and rising sales generate further consumer acceptance.

In January 2015, the PBOC invited eight private entities to start “preparatory work” for a pilot personal credit scoring scheme, each with its own data advantage. The choice largely reflected the growing importance of big data analytics in the digital era. Zhima (Sesame) Credit, Tencent Credit Information, and Qianhai Credit Information all have their own internal data sources coming from e-commerce, social media, payment and other financial services, besides the profuse use of a significant amount of external data. For instance, Qianhai Credit’s data include those from Ping An Group’s subsidiaries involved in various traditional financial services such as insurance, banking, trusts, securities, and funds, and fintech services such as Lufax P2P platform, OneConnect Financial Technology, and 1qianbao e-wallet.

Koala Credit and Huadao Credit, each initiated by multiple shareholders, have multidimensional data sources ranging from payment, financing, wealth management to other financial services. For instance, Shenzhen InfoTech Technologies, the controlling shareholder of Huadao Credit, has worked with more than 300 banks on financial big data. As established credit reporting agencies, China Chengxin Credit and Pengyuan Credit have accumulated a large amount of data, most of which come from financial and third-party institutions including the public sector. China Zhicheng (IntelliCredit) has cooperated with more than 1,000 P2P platforms and has a large amount of credit blacklist data.

The pilot scheme eventually failed to gain PBOC approval, for several possible reasons.⁸ First, tech companies' "credit scores" were widely used as a proxy for trustworthiness, far beyond the scope of credit scoring. Second, there was a perceived lack of independence, as well as market distortions and conflict of interests claimed to have arisen from some of the agencies' close links to in-house businesses. Third, each pilot agency had its own data sources neatly tied to a relatively closed business and data loop, resulting in segmented information and conflicting scores given to the same individuals by different rating firms depending on their use of specific products or services under each firm's coverage. Each firm had distinct user bases: over 80% of Alibaba's MYbank customer base are MSMEs that had never obtained financial services from the banking system, while Tencent's WeBank targeted individual consumers; Qianhai and Koala covered a considerable number of internet financing users. Following the PBOC's decision, these eight pilot institutions have subsequently withdrawn from personal credit reporting, some of them began to focus on data services or on social and commercial rating schemes.

A landmark event in China's credit reporting history was the incorporation of Baihang Credit, China's first and only company with a personal credit business license, on March 19, 2018. The National Internet Finance Association of China (NIFA) and the eight entities in the pilot scheme became its shareholders. Besides basic personal information and bank loan records, Baihang Credit analyze mobile phone and online shopping data, credit records from online lending platforms, traffic violations, legal disputes, and overdue utilities bills. However, Baihang Credit faces significant challenges in convincing its main shareholders to transfer relevant proprietary data. This request involves serious user privacy issues, as any data transfer unauthorized by the user is deemed illegal. Intended to be an independent firm providing credit scoring services, Baihang Credit engages in personal and corporate credit reporting and database management, under PBOC supervision and guidance. In this new "government + market" structure, the CRC focuses on basic services while Baihang Credit provides value-added services for personal credit inquiries. In July 2020, the CRC and Baihang Credit signed a strategic cooperation agreement to jointly pursue research on credit reporting strategies, businesses and technology.

⁸ See, for example, Lian, Bian, Su and Cao (2017) and Chorzempa (2018).

3.1.5 The Case of Zhima Credit⁹

Zhima Credit, also known as Sesame Credit, is a typical case of using big data to enhance social and commercial trust building in China. In January 2015, Ant Group's Sesame Credit started as a trial run of a personal "credit-scoring system", with the initial purpose of improving credit access for millions of less privileged borrowers. Over time, the focus of these ratings has shifted away from providing FICO-type scores for lenders, toward offering merchants and service providers a reliable tool for assessing customer trustworthiness, and enabling greater consumer access to a wide range of services. Zhima Credit hence stopped cooperating with financial institutions. Based on big data analytics using both online and offline behavioral data, Zhima Credit generates "trust scores" for consumers and small business owners. Zhima scores are computed based on data ranging from credit history to social networking. Besides data sources within Alibaba and Ant Groups, most data come from hundreds of data partnerships. The process goes from user authorization, data collection and processing to model estimation and the production of final scores.

From early on, Zhima Credit focused on commercial trust building and on serving the real economy, having already provided trust services in nearly 40 industries. Zhima Credit focuses on several dimensions. First, Alibaba's internal ratings and Zhima scores facilitated transactions on various commercial and financial platforms. Internally, it afforded important risk control services to, e.g. Ant Credit Pay (Huabei), a digital advance payment product, and Ant Cash Now (Jiebei) consumer loan service. In July 2016, Zhima Credit launched its Lingzhi system that provides insights into the credit risk profiles of numerous MSMEs. In August, three small loan platforms including Zhaolian Good Lending and Lai Instalment launched loans with floating interest rates linked to Zhima scores, making it easier for borrower with high Zhima scores to obtain cheaper funds.

Second, Zhima trust scores have served as a catalyst of China's growing platform and sharing economy. A core issue is trust: a qualified renter may be unable or unwilling to pay a substantial sum of deposit before leasing for fear of potential loss, but the lessor is reluctant to provide services without such a deposit. In July 2018, Zhima Credit advanced the concept of a trust-based "new rental economy", providing a solution to this dilemma. Specifically, Zhima trust scores were used to enable trades by waiving the deposit based on

⁹ Distinct from other credit reporting agencies which name themselves "Credit Scoring" firms, a more appropriate translation of "芝麻信用" should be "Zhima Trust" rather than "Zhima Credit", reflecting its focus on helping build trust among people and facilitate commercial transactions.

the assessment of a user's trustworthiness. Since its inception, Zhima Credit has pioneered numerous deposit-free rentals in cooperation with short-term rental platforms (cars, bikes, housing, books, power banks, clothes, etc.). For example, in January 2015, in collaboration with Shenzhou Car Rental, Zhima Credit began to provide the world's first purely deposit-free car rental service. It collaborated with Fliggy Travel to provide deposit-free hotel reservations for high-score users. In July, it joined forces with Yongan Bicycles to provide deposit-free bike rentals. In 2016, Zhima Credit cooperated with McDonald's to provide deposit-free umbrella and power bank rentals; with 1zu.com, a real estate management platform, to allow high-score users to rent a house without deposit; and with Shanghai Library to waive residents with high Zhima scores the deposit needed to apply for library cards and borrow books. Zhima scores have played a key role in promoting new sharing platforms in China.

A third dimension is trust building and fraud detection, prevention, and deterrence. In July 2015, Zhima Credit obtained the Supreme People's Court's approval to provide real-time updates on its list of "scroungers", i.e. debtors with arrears long overdue. They would then receive payment reminders from Alipay, and be penalized with reduced Zhima scores leading to curtailed access, according to law, to services provided by Zhima Credit partners. Under pressure, over 15,000 "scroungers" paid off their debt within the first two months, of whom over one-third had evaded execution for more than two years. This turned out to be more effective than the travel or borrowing restrictions imposed by the Court. In 2016, Zhima Credit and Guizhou Traffic Bureau jointly launched a system that would evaluate and score the drivers of "important vehicles"; and Zhima Credit provided ID authentication services to dating platforms Huatian and Jiayuan.com to prevent fraudulent behavior.

Zhima Credit has also provided other types of trust services. In June 2015, along with Ali Travel, Zhima Credit allowed faster and more convenient visa applications to Singapore for users with a score of 700 or above. In November, 2018, Chinese nationals who had Zhima scores over 750 could submit a Zhima report as proof of financial standing in their visa application for Canada, and the Canadian government stated that "the Zhima report is more comprehensive than the traditional bank statements and deposit certificates typically received in support of visitor visa applications". In 2016, Zhima Credit and Guangzhou Women's and Children's Medical Center jointly launched the "Pay for the first consultation and pay later" service, saving more than 60% of a patient's visit time. In May 2017, China Unicom users with Zhima scores of 650 or more would not need to make

advance payments for certain mobile phone service packages, allowing them to make calls before paying phone charges.

Many other tech companies have followed suit since then. Tencent Credit began testing its WeChat Pay Points in mid-2017, which were officially launched in early 2019. The system leverages artificial intelligence and rates a user's trustworthiness based on her spending behavior and personal connections, among other factors. With over 100 million WeChat Pay points users (as of January 2020) and services covering more than 30 industries, already 80% of users have started to enjoy deposit-free services. JD.com, the second largest e-commerce platform in China, runs a consumer risk assessment system launched in 2015, in a joint venture with ZestFinance, a US firm that uses machine-learning in big data analysis to produce more accurate credit scores.

3.2 Data Privacy and Algorithms Fairness

Fundamental to any successful trust building mechanism and a credible credit reporting system is a faithful record of past behavior based on reliable data, much of it being of private nature. In a digital world, individuals are not only consumers of information, but also producers of large volume of often highly sensitive data. Almost everything we do leaves digital or physical trails that generates data, and increasingly, it often goes beyond our ability to control the way such data is collected, used, and shared. Privacy, acceptable use, and data security assume an increasing role in the new era of big data, posing complex challenges for the credit reporting industry.

As the scale and scope of data collection continues to expand, concerns have been growing with data privacy and transparency (e.g. how data is collected and used), besides the obvious issues of data protection and cybersecurity. Acquisti, Taylor, and Wagman (2016) survey the privacy literature. They find that in a digital economy, consumers' ability to make informed decisions about their privacy is severely hindered because "consumers are often in a position of imperfect or asymmetric information regarding when their data is collected, for what purposes, and with what consequences". Yet personal data has value, and data sharing may reduce market frictions and facilitate transactions. The increasing amount of information can yield significant benefits to both the data subjects and data collectors. Luohan Academy (2021) finds that disallowing the use of personal data for online marketing hurts the consumers as they lose valuable personalized sales recommendations. As a result, with little knowledge of customers' revealed preferences, buyer-seller matches and online purchases drop significantly. Moreover, the

well-known “privacy paradox” is not puzzling: while people genuinely care about their privacy, they are rational and willing to provide sensitive information in exchange for expected gains from access to online services.

Data privacy laws, rules, and principles vary widely across different jurisdictions. A major task for policymakers is to strike the right balance between reaping the benefits of fluid, expanding information flows, and reducing risks with privacy violations and data breaches as well as algorithmic unfairness in credit reporting systems. In the United States, concerns with privacy led to the 1970 Fair Credit Reporting Act (FCRA), which established strict limitations on the access to and use of data, as well as consumer rights enshrined to ensure the transparency and accuracy of the data used in credit reports.¹⁰ The FCRA, substantially revised in 1996, 2003, and 2018, provided a framework aimed at ensuring fairness, accuracy, and privacy of the personal information used for credit reporting. The California Consumer Privacy Act, enacted in 2018, focused on consumer rights relating to the access to, deletion of, and sharing of personal information that is collected by businesses.

In Europe, the European Union’s General Data Protection Regulation (GDPR), effective in May 2018, set a benchmark for privacy and data protection in the information age. It regulates the collection, storage, and processing of personal data, i.e. any data that linked to a specific person including direct (e.g. name, national ID number) and indirect personal identifiers (e.g. phone numbers, IP addresses, or photos). In most cases, the processing of personal data is allowed only if the data subject has given its consent, and automatic decision-making processes with an impact on individuals, including profiling, as well as the processing of extremely sensitive data (e.g. biometric data) require “explicit” consent. The data minimization principle limits personal data collection, storage, and usage to data that are relevant, adequate, and necessary for carrying out the purpose for which the data are processed. The storage limitation principle restricts the duration of data storage to a specified (necessary) period.

However, big data analysis in trust reporting relies on large amounts of data, it often starts to work with certain data without a specific purpose a priori, this is inconsistent with the principle of data minimization. In big data analysis, very often new hypotheses for testing are introduced after data is collected. However, the data subjects might have given their initial consent for data collection for a different purpose. A safe way to conduct such analysis in a way that conforms to GDPR requirements is to process only anonymous data, i.e. data that does not include direct personal identifiers. But

¹⁰ See Hurley and Adebayo (2017) for a detailed account.

this renders big data analysis less valuable for each individual. Limited data storage shortens data span, limiting a models' predictive ability and potentially hurting the thin-credit-file customers in the fat tail, who may need it most to build up their credit record over time.

In China, the 2013 "Regulations on the Administration of the Credit Reporting Industry" establish the basic rights of users, namely their right to consent (data collection, inquiry, and use must have user consent for intended uses); right to access to data and to delete bad information (data cannot be kept for more than five years); right to inquire, dissent, and appeal; and right to data security. The Regulations stipulate that credit reporting agencies "shall not collect information about the income, deposits, negotiable securities, commercial insurance, real property or taxes of individuals, unless they have expressly informed the individuals concerned of the possible adverse consequences of providing such information and have obtained their written consent".

In reality, however, relevant legislation is still lacking, and grey areas exist due to rapid changes in credit reporting. Many organizations collect private personal information illegally under the guise of big data credit reporting, and their use of the data is often unethical or even criminal. More efforts are made by both regulators and tech companies to establish a solid framework that supports sound data use and credit reporting practices. Furthermore, private initiatives aimed at beefing up information security are important. In April 2016, Zhima Credit became the first credit reporting agency in China to obtain ISO/IEC 27001:2013 certification. The agency applies desensitization technologies broadly to protect user privacy even from in-house engineers, with Zhima trust scores being produced for individual users through automated algorithms that preempts unwanted human interference.

Automated ratings became more widely adopted in the early 1990s, breaking the dominance of subjective decision-making by loan officers and experts. With this, important social issues also emerged in the use of scoring as a forecasting tool, including algorithmic fairness and accountability. In response, US Equal Credit Opportunity Act enacted in 1974 prohibits lenders from discriminating against loan applicants on the basis of race, gender, national origin, religion, or marital status. China's 2013 Regulations stipulate that "credit reporting agencies are prohibited from collecting information about religious beliefs, genes, fingerprints, blood types, diseases or medical history of individuals, or any other personal information prohibited by laws and administrative regulations".

As automated big data analytics increasingly replace human decision-making in the digital age, questions arise concerning algorithmic ethics,

fairness and transparency. While it is illegal or unethical to use individuals' certain traits or attributes such as gender, race, age, or religion, there are growing controversies concerning the use of surrogate or correlated variables which allow a scoring system to discriminate in these dimensions. Algorithms should be fair and transparent, they should not be seen as promoting discrimination in hiring, lending, or judicial decisions. An individual should not be denied credit or jobs for reasons unknown to her, and algorithms owners should be held accountable for automated decisions.

Automation, especially allied with big data analytics, was expected to increase efficiency and prevent the more grievous forms of discrimination. Nevertheless, as Hurley and Adebayo (2017) suggest, it might have actually further aggravated discrimination issues. Kleinberg, Ludwig, Mullainathan, and Rambachan (2018) point out that the "growing use of algorithms in social and economic life has raised a concern: that they may inadvertently discriminate against certain groups". Nevertheless, using nationally representative data on college students, they find that the inclusion of race in application actually improves efficiency (predicted GPAs of admitted students) as well as equity (fraction of admitted students who are black), and "the strategy of blinding the algorithm to race inadvertently detracts from fairness". Variables like age, gender, and race can be included in algorithm-based big data analytics along with equity preferences to better promote fairness.

4 Conclusion

Trust is a fundamental element of the platform-based digital economy, and efficient trust building and credit reporting systems are of crucial importance to the provision of many different types of commercial and financial services, especially in areas that information asymmetry is most severe, e.g. e-commerce, insurance, and credit intermediation. While the Chinese economy has experienced a "growth miracle" in the last four decades, its credit reporting system still lags far behind that in many advanced economies, hindering further high-quality growth.

China's credit reporting system has come a long way, transforming from a rudimentary government-led initiative into a mixed system, where both the official Credit Reference Center and a highly competitive private credit reporting market play an important role. While credit reporting coverage has grown over time, the services of Baihang Credit, the new credit bureau, are seen by some as a promising way to provide more effective solutions to satisfy

the growing needs for accurate and timely credit reporting in a market with segmented information and dazzling technological advances.

Digital technologies have helped usher us to an age of explosive information. Big data analytics, based on the widening range of data sources, growing storage and computing capabilities, and combined with AI-enabled automated algorithms and decision-making, have led to exciting advances in social and commercial trust building as well as credit reporting, especially in China where there is a high perceived trust deficit. An efficient trust reporting system promotes economic stability by helping firms better price and control risks. Nevertheless, policymakers and tech companies face the intriguing balancing act between reaping the benefits from innovative, digital trust building, and addressing the important issues concerning consumer protection, privacy, data security and algorithmic fairness. Already, internal rating schemes and trust scores built on new technologies and provided by some tech companies have opened up new possibilities, enhancing efficiency and supporting, for example, the rise of e-commerce, digital platforms and the sharing economy.

Acknowledgements We thank Wei Tan and Yi Zhou for helpful conversations, and Yuchi Zheng for assistance with graphs. The views expressed here belong to the author and do not necessarily reflect those of Luohan Academy, Alibaba Group, Ant Group, or China International Capital Corporation Limited.

References

- Acquisti, Alessandro, Curtis Taylor and Liad Wagman. 2016. The economics of privacy. *Journal of Economic Literature* 54 (2): 442–492.
- Akerlof, George A. 1970. The market for “Lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84 (3): 488–500.
- Backer, Larry Catá. 2019. China’s social credit system: Data-driven governance for a ‘New era’. *Current History*, 209–214.
- Barron, John M., and Michael E. Staten. 2003. The value of comprehensive credit reports: Lessons from the U.S. Experience. In *Credit Reporting Systems and the International Economy*, ed. Margaret Miller, pp. 273–310. MIT Press.
- Berg, Tobias, Valentin Burg, Ana Gombović and Manju Puri. 2019. On the rise of FinTechs: Credit scoring using digital footprints. *Review of Financial Studies*.
- Berger, Allan, Nathan Miller, Mitchell Petersen, Raghuram Rajan and Jeremy Stein. 2005. Does function follow organizational form? Evidence from the lending practices of large and small banks. *Journal of Financial Economics* 76: 237–269.

- Chorzempa, Martin. 2018. China needs better credit data to help consumers. *Policy Brief 18-1*, Peterson Institute for International Economics.
- Chorzempa, Martin, Paul Triolo and Samm Sacks. 2018. China's social credit system: A mark of progress or a threat to privacy? *Policy Brief 18-14*, Peterson Institute for International Economics.
- Diebold, Francis X. 2019. *On the origin(s) and development of 'Big data': The phenomenon, the term, and the discipline*. Working Paper, University of Pennsylvania.
- Dietrich, David, Barry Heller and Beibei Yang. 2015. Data science & big data analytics: Discovering, analyzing, visualizing and presenting data, EMC education services, John Wiley & Sons, Indianapolis.
- Durand, D. 1941. Risk elements in consumer installment financing. National Bureau of Economic Research, New York.
- Einav, Liran, Mark Jenkins and Jonathan Levin. 2013. The impact of credit scoring on consumer lending. *RAND Journal of Economics* 44 (2): 249–274.
- Gambacorta, Leonardo, Yiping Huang, Han Qiu and Jingyi Wang. 2019. How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm. BIS Working Papers No 834.
- Huang, Zhuo, Yang Lei and Shihan Shen. 2016. China's personal credit reporting system in the internet finance era: challenges and opportunities. *China Economic Journal* 9 (3): 288–303.
- Hurley, Mikella and Julius Adebayo. 2017. Credit scoring in the era of big data. *Yale Journal of Law and Technology* 18 (1): 148–216.
- Liberti, José María, and Mitchell A. Petersen. 2019. Information: Hard and soft. *The Review of Corporate Finance Studies* 8 (1): 1–41.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan and Ashesh Rambachan. 2018. Algorithmic Fairness. *AEA Papers and Proceedings* 108: 22–27.
- Laney, Doug. 2001. 3-d data management: Controlling data volume, velocity and variety. META Group Research Note, February.
- Lazer, David, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. The parable of Google flu: Traps in big data analysis. *Science* 343 (6176): 1203–1205.
- Lian, Wei, Hui Bian, Xianghui Su and Pengcheng Cao. 2017. Ant financial: From Alipay to a new financial ecosystem, Renmin University of China Press.
- Liu, Chuncheng. 2019. Multiple social credit systems in China. *Economic Sociology: The European Electronic Newsletter* 21 (1).
- Louzada, F., A. Ara, and G. B. Fernandes. 2016. Classification methods applied to credit scoring: Systematic review and overall comparison. *Surveys in Operations Research and Management Science* 21 (2): 117–134.
- Luohan Academy. 2021. Understanding big data: Data calculus in the digital era, report.
- Miller, Margaret J. 2003. Credit reporting systems around the Globe: the state of the art in public credit registries and private credit reporting firms. In *Credit Reporting Systems and the International Economy*, ed. Margaret J. Miller, MIT Press.

- Munkhdalai, Lkhagvadorj, Tsendsuren Munkhdalai, Oyun-Erdene Namsrai, Jong Yun Lee, and Keun Ho Ryu. 2019. An empirical comparison of machine-learning methods on bank client credit assessments. *Sustainability* 11: 699.
- Ohlberg, Mareike, Shazeda Ahmed and Bertram Lang. 2017. Central planning, local experiments: The complex implementation of China's social credit system. MERICS China Monitor, Mercator Institute for China Studies.
- Olegario, Rowena. 2003. Credit reporting agencies: A historical perspective. In *Credit Reporting Systems and the International Economy*, ed. Margaret J. Miller, MIT Press.
- Síthigh, Daithí Mac and Mathias Siems. 2019. *The Chinese social credit system: A model for other countries?* EUI Working Paper LAW 2019/01.
- Stein, Jeremy C. 2002. Information production and capital allocation: Decentralized versus hierarchical firms. *Journal of Finance* LVII (5): 1891–1921.
- Stiglitz, Joseph. 2000. The contributions of the economics of information to twentieth century economics. *Quarterly Journal of Economics* 115 (4): 1441–1478.
- Tirole, Jean. 2020. Digital dystopia. Working Paper, Toulouse School of Economics.
- Thomas, Lyn C. 2000. A survey of credit and behavioural scoring: Forecasting financial risk of lending to consumers. *International Journal of Forecasting* 16: 149–172.
- World Bank. 2019a. Doing Business 2020. Washington: World Bank.
- World Bank. 2019b. Credit Reporting Knowledge Guide 2019. Washington: World Bank.
- Zeckhauser, Richard, and Paul Resnick. 2002. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. *The Economics of the Internet and E-Commerce*, Ed. M. R. Baye, Elsevier Science, 127–157.
- Zhang, Yingjie, Yabin He, Qian Liu, Liang Chen and Bing Liu. 2019. 中国社会信用体系建设与改革40年,” 中国诚信研究院.



Digital Identity in India

Bhagwan Chowdhry, Amit Goyal, and Syed Anas Ahmed

The World Bank (2016) identifies three main goals of any identification system:

1. Inclusion and access to essential services such as health care and education, electoral rights, financial services, and social safety net programs.
2. Effective and efficient administration of public services, transparent policy decisions and improved governance—particularly to reduce duplication and waste.
3. More accurate measure of development progress in areas such as reduction in maternal and infant mortality.

¹ The ID2020 Alliance (<https://id2020.org/>) is a worldwide initiative to create digital ID through a multi-stakeholder partnership, ensuring digital ID is responsibly implemented and widely accessible.

B. Chowdhry (✉) · S. A. Ahmed
Indian School of Business, Hyderabad, India
e-mail: bhagwan@isb.edu

S. A. Ahmed
e-mail: Syed_Ahmed@isb.edu

A. Goyal
Swiss Finance Institute, University of Lausanne, Lausanne, Switzerland
e-mail: amit.goyal@unil.ch

With these goals in mind India embarked on a massive project in 2009 to generate a digital identity for more than one billion residents of the country.¹ Nandan Nilekani, an IT technologist who co-founded India's most iconic IT company, Infosys, was asked to oversee the creation and implementation of this project—creating a new cabinet minister level position—which was later given a name, Aadhaar, which in Hindi means “the foundation” (Parker 2011).

The role of an identity is to assert that you are who you say you are. To be more specific, the assertion links a physical body in the present moment to the physical body of the same person at another moment in time. So, for example if I go to a university and claim that I am Bhagwan Chowdhry, I am asserting that the person who is making this claim is the same Bhagwan Chowdhry who went to a school called St. Paul's in the city of Ajmer and received 98% marks in his high school examinations. This requires many previous steps of identity assertion and verification. When Bhagwan went to his high school, he also had to assert his identity when he took the high school exams in which he claims to have scored 98% marks.

Many identity protocols have existed, most beginning in the twentieth century. The most common is a photo identity in which a name of a person is linked to a picture of the person at the very minimum and is certified by some institution—a government, a school, a university, or a financial institution. Other data—a signature, a date of birth, a residential address, name of a family member (e.g., a parent or a spouse)—could be added to the picture identity. Forged and fake identities are ubiquitous and create nuisance, inefficiencies, and corruption at many levels.

A second widely used identification protocol is to link a unique number to each person. The ten-digit social security number in the U.S. is the most prominent example. The uniqueness of the number avoids the problem associated with names which may not be unique and may have variant spellings. If many events or transactions can be linked to the same unique number, it creates a linked history which can be used (or abused) for various purposes. The algorithm used to create and assign the ten-digit social security number is neither random nor particularly private, nor is it easy to ensure that no one is able to obtain more than Social Security Number; it was designed decades ago when privacy issues were not salient.

1 What Is Aadhaar?

Aadhaar strived to create an identity that would be unique, safe, portable, and extremely difficult to forge. Aadhaar provides a unique, randomized 12-digit number (that means 1 trillion numbers are available) to be assigned to over one billion residents in India. Since only one in a thousand numbers is used for each person, randomization makes it difficult to guess which number belongs to whom (UIDAI 2017).

This Aadhaar number is then bound to unique biometric markers for each individual. The biometric markers are fingerprint scans of ten fingers, retina scans of two eyes and a photograph. The biometrics are stored digitally (OECD 2018).

2 Technological Feat

A major challenge in generating and assigning a unique number to each individual's biometric markers is ensuring that no individual gets more than one Aadhaar number. This means that when the digital biometric data is collected and stored in computer servers, the system needs to check that the same data has not already been used to obtain another Aadhaar number already. As the number of registrations grow, potentially reaching nearly one billion, this involves checking the biometric data against biometric data for all individuals already registered and stored in the system. This is a non-trivial computation problem. It is worth noticing comparing digitized biometric data is not simply comparing numbers but rather looking for patterns in the digitized data to determine if they came from the same individual. This was a sophisticated technological feat achieved by Unique Identification Authority of India (UIDAI) engineers (Varma 2010; Gunipati 2017).

Another challenge was enrollment and capturing the biometric data. First it requires equipment—fingerprint readers and iris scanner cameras. The equipment had to be acquired and deployed all over the country. UIDAI, in addition to government centers, also incentivized private players to obtain the equipment, enroll the residents by paying a fee to the private operators for each successful enrollment and imposing a small penalty for any errors in capturing the biometric data. This was a huge undertaking but a relatively successful one. In just five years after the rollout, nearly 600 million residents were enrolled and by the end of 2019, nearly twice that number has already been enrolled.

This is nothing short of a miracle in a country where a large population is not highly educated, lives in villages and does not even bother to obtain a traditional identity such as a birth certificate.

3 Incentives and Coercion

Financial exclusion for a substantive fraction of the poor population in developing countries such as India is a serious concern. It has been recognized in the theoretical literature (Aghion and Bolton 1997; Banerjee and Newman 1993) that financial inclusion can help people transform their production and employment activities to enable them to exit poverty. Sen (2000) suggests that the very provision of services such as banking and credit enables generation of capabilities boosting economic growth. Reserve Bank of India's Committee on Financial Inclusion notes in its report (Rangarajan 2008) that financial inclusion to hitherto excluded segments of the population is critical to sustain and accelerate growth.

Many scholars and policy makers have proposed and attempted various ideas for financial inclusion. One prominent example of improved financial inclusion is through microfinance institutions such as the Grameen bank in Bangladesh. For these experiments, lack of identity is not a stumbling block as social contracts and collaterals are utilized in facilitating transactions and access to credit. However, the Grameen experiment also indicates the limits of the experiment and suggests that government intervention may be necessary to achieve bigger scale. Thus, the policy discourse has shifted from nudge to actively promoting the banking sector for financial inclusion with some arguing for creating incentives. For instance, in 2009, Financial Access at Birth (FAB) was proposed in which an incentive of \$100 dollar electronic bank deposit for every child born would be created to enroll the child and the parents in an electronic identity protocol linked to a savings bank account. Of course, this required substantial infusion of funds.²

Unlike the informal systems, formal systems of banking depend on codified information about customers. Digital identity can facilitate opening bank accounts thereby granting access to the formal financial sector. Note, however, that solving the demand side problem does not necessarily obviate the need for solving the supply side issue. If a business case for inclusive banking cannot be made for the banks, then digital initiatives may not achieve the

² <http://financialaccessatbirth.org/>.

desired scales. Thus, the last mile problem remains even with the advent of digital identity (Sriram 2014).

Recent theoretical literature (Besley and Persson 2009, 2010) highlights the importance of state capacity for economic development. Even though billions of dollars are spent worldwide on anti-poverty programs, one constraint in effective implementation is the lack of secure payment infrastructure to make payments to intended beneficiaries. Digital identity can improve access to government services as it makes it easier for the government to both accurately identify intended beneficiaries and transfer benefits directly to their bank (if digital identity also helps beneficiaries in opening bank accounts).

UIDAI by linking many Direct Benefit Transfers (DBT), such as food subsidy programs that are quite prevalent in India, to Aadhaar played the dual role of creating incentives and effective implementation of the DBT programs (Pande 2014). The prime minister of India Narendra Modi in 2014 implemented a program—called PMJDY—incentivizing opening and linking of bank accounts to Aadhaar by offering free life insurance to those who enrolled.³

It became clear, however, that incentives alone would not be enough to achieve near universal enrollment. For a large period until the Supreme Court of India made a ruling in 2018, it was not clear if Aadhaar was a mandatory requirement or voluntary. Many people felt it was mandatory and many private and public service providers acted as if it was. This confusion and resulting indirect coercion helped in achieving massive enrollment much to the chagrin of many civil society members of the country.

4 Costs and Benefits

Some have criticized the Aadhaar program by complaining that proper cost-benefit analysis was not performed before embarking on this ambitious program. This is correct in the sense that no formal analysis was done. This was a political decision by the government of previous Prime Minister Manmohan Singh of Congress Party who judged, perhaps informally, that at a very broad level this was an investment in infrastructure that would lead to a substantial increase in social welfare. When Narendra Modi of the opposition party BJP became the Prime Minister in 2014, he too was convinced by

³ <https://pmjdy.gov.in/>.

Nandan Nilekani's persuasion that Aadhaar would provide substantial benefits to the new India that was about to embrace digital technology in a big way (Aiyar 2017). It is thus incorrect to assert that no cost–benefit analysis was performed in making a decision to introduce and then to continue with the implementation of Aadhaar. Such informal calculations are not uncommon for many important political decisions of great significance—formal analyses are often a facade to justify decisions that have already been made.

The budget for Aadhaar was \$1.5 billion and so the direct cost per identity created was less than \$2. It is not unreasonable to estimate that the lifetime benefits of a unique ID might be many folds than that number. What are some of these benefits? Direct Benefits Transfers in India amount to approximately 3–4% of yearly GDP. Estimates of leakages vary but even with increasing efficiencies over the years, the leakages amount to tens of billions of dollars every year. Other positive externalities created by a fast, efficient identification protocol that could substantially reduce the costs of obtaining financial and other services probably dwarf the direct costs of implementing Aadhaar.

5 Indirect Costs

Critics of Aadhaar have argued that these benefits notwithstanding, the design of Aadhaar is flawed because it creates other injustices and inequities that were not carefully thought of in a comprehensive manner (Khera 2019).

6 Linking Digital Identity with Public Distribution System

Proponents of Aadhaar have often cited the benefits of linking a digital identity with public distribution system. However, there are some valid reasons to doubt the effectiveness of linking public distribution programs to digital identities. First, most-excluded beneficiaries might also be the same as those who are unable to obtain digital identities (Khera 2011) thus hurting the same segment of population that the program is designed to help the most. Second, implementation snafus might lead to failure of the undertaking limiting the program's effectiveness (Kremer 1993). Third, vested interests whose rents are threatened may subvert the implementation of the program (Krusell and Ríos-Rull 1996). Relatedly, reducing corruption could displace it

to other sectors (Yang 2008) or could even dampen the incentives for officials to implement these programs (Leff 1964; Niehaus and Sukhtankar 2013).

Jean Drèze, a Belgian born Indian economist and activist who has spent his lifetime working in India on development issues such as poverty, gender inequality, and financial exclusion has argued that a slight improvement in the identity system used in India for decades, known as the ration card which identifies and lists family members in a household, to a smart card (Bangladesh, for example, has implemented these⁴) would have led to many of the benefits associated with leakage of food subsidy transfers touted by proponents of Aadhaar.

Drèze worries that requiring Aadhaar to claim food subsidies, in fact, creates many difficulties for the very poor and marginalized sections of the society. This is because Aadhaar verification requires technological equipment, to scan fingerprint and irises, and ubiquitous high-speed internet that works all the time. Cases of people being denied their food rations because their Aadhaar identities could not be verified because of equipment failure, poor internet connectivity, or other behavioral difficulties highlighted these concerns (Ranjan 2018). Proponents of Aadhaar claim that such failures are minuscule compared to the cases of flawless identification but that does not pacify social activists such as Drèze who worry that it is the very poor who are severely hurt by these failures and their numbers being small is little consolation for those affected.

Even beyond technical failures in ensuring identification is the risk that formalizing identification might lead to a stronger divide between insiders and outsiders. Due to reduced incentives/opportunities of insiders to game the system, undocumented people might be further marginalized (the golden goose effect; Niehaus and Sukhtankar 2013). There is some evidence that Dominican Republic's 2006 experiment led to increased statelessness (Gelb and Clark 2013).

7 In the Pocket or in the Cloud?

For Nilekani, the fact Aadhaar is just a number and that biometric data could be stored and verified in the cloud, anywhere in the country, any time of the day, was the disruptive innovation—he describes this realization as an epiphany. Aadhaar obviates the need to carry anything in your pocket, it does not get lost, and cannot be forged, this was utterly remarkable.

⁴ https://en.wikipedia.org/wiki/Bangladeshi_National_Identity_Card.

How does identity verification work in Aadhaar? Suppose I go to a bank to open an account and say, “I am Bhagwan Chowdhry, my Aadhaar number is 1234 4321 6789 and I would like to open a bank account.” The bank scans my fingerprints and sends a query to UIDAI’s servers asking, “Can you verify that the person whose fingerprint we have sent you a scan of, is indeed linked to the Aadhaar number 1234 4321 6789?” Notice that verifying this information requires a trivial computation that can be performed nearly instantaneously. The servers send back an answer in the form of a “Yes” or a “No.” No other information is transferred from the UIDAI servers to the bank and the UIDAI servers do not keep or store that this query was made by a particular bank to verify Bhagwan’s identity.

True, this required a huge initial infrastructure costs but as a public infrastructure project, the per capita costs were trivial. Nilekani, with other technology optimists, further imagine that internet access will become very inexpensive and ubiquitous and financial inclusion will proceed rapidly in the country as the infrastructure develops (Nilekani 2018). Their focus is on India in the coming decades whereas Drèze and other critics are concerned about the damage to marginalized populations now.

8 Data and Privacy

All centralized systems cause many people to worry about another important issue—data privacy. These issues arise from linking additional demographic and financial data to the identity. This is particularly problematic because biometric data, unlike a password, cannot be changed if someone were to steal a person’s biometric data.

In addition to the photograph and biometric data, Aadhaar also lists the address where the individual resides. Notice that this additional data is not absolutely necessary for creating the unique ID. The fact that this data exists in central servers of the government raises the fears that government, and other nefarious actors, may engage in surveillance activities because they know where each individual resides. Servers could be hacked and the data could be compromised.⁵

Recent advances in computer science (Agarwal, Banerjee, and Sharma 2017) provide a solution to the hacking problem. This involves not storing

⁵ At the same time, some linkages can be beneficial as they reduce administrative overhead. For example, Argentina linked tax, real property, and social service data to reduce fraud and tax evasion. The estimated benefits of this program at \$104 million far outweighed the implementation costs of \$10 million (Gelb and Clark 2013).

the unencrypted data at all in any centralized servers but instead encrypt the data and distribute partial data across many servers geographically located at many independent institutions, such as universities, civil society organizations, etc., such that if any of these servers is hacked, the hacker will not get any useful data at all but when a legitimate verification query arrives, encrypted data from various servers is recombined and verified using zero knowledge proofs protocols (Chaum et al. 2017).

But this does not eliminate the data privacy issues (Khera 2018). This is because the identity data linked with other activities and transactions is also collected by other service providers such as banks. A printed Aadhaar card with a photograph, the Aadhaar number, and address has de facto become a photo identity that is being used as identity verification at airports, hotels, banks, and other service providers. If these organizations begin collecting data and store it electronically in databases, those databases are also subject to hacks making sensitive data vulnerable (Economist 2018a).

To minimize some of these issues UIDAI introduced the concept of a Virtual ID or VID which is a 16-digit number temporarily linked to the permanent 12-digit Aadhaar number. A VID can be used for authentication just like the Aadhaar number and a service provider cannot derive the original 12-digit Aadhaar number from the 16-digit VID.

To further address the data security issues, the Supreme Court of India in its 2018 judgment ruled that Aadhaar cannot be made mandatory for provision of most services except for filing tax returns and an associated identity known as the PAN card. Linking of all other services such as banking would require that parliament pass a separate law providing such permission in a limited and controlled fashion. The Supreme Court also directed the government to debate and pass an explicit data protection and privacy law by the parliament. The parliament proposed a new data privacy law in December 2019 that severely restricts private corporations from using consumers private data but it exempted the government itself from such restrictions if it is deemed necessary for “national security” (Wimmer and Maldoff 2019). The tradeoffs for the very poor, who are essentially excluded both socially and financially, and the relatively well-off who worry about their privacy a bit more, are starkly different. The regulatory framework, however, is uniform for all people. The issues and the solutions that evolve, in the context of India, often represent a *compromise* of “dignity and dignified living” (Singh 2019).

9 Survey Evidence from India

Dalberg conducted a pulse survey with 147,868 households in 28 states and union territories, and an in-depth survey with 19,209 households in 16 states and 1 union territory. Its findings from its State of Aadhaar report (Dalberg 2019) are reproduced below

1. Aadhaar is becoming ubiquitous in India

95% of adults have Aadhaar, and on an average use it once a month.
75% of children have Aadhaar.

2. A notable minority still does not have the ID

90% of residents in Assam and **61% of residents** in Meghalaya do not have Aadhaar.

30% of homeless, and 27% of third-gender residents do not have Aadhaar.

8% of people do not have Aadhaar—or an estimated 102 million people, 75 million of whom are children.

3. Updating is the hardest part of the Aadhaar process

33% of people who tried to update found the process difficult; one in five did not succeed.

4% of people currently have errors in the information on their Aadhaar card.

15% of people have an error in their linked mobile phone number; an additional **39%** have not linked a number at all.

4. Aadhaar has supported inclusion

49% of people used Aadhaar to access one or more services for the very first time (e.g., ration, MGNREGS, social pensions, SIM cards, and/or bank accounts).

For 8% of people, Aadhaar was their first ever ID.

5. For many residents Aadhaar has improved service delivery

80% of beneficiaries feel Aadhaar has made PDS rations, MGNREGS and social pensions more reliable.

Using Aadhaar, residents were 40% more likely to obtain a new SIM card within one day, compared to using other forms of ID.

6. Problems with Aadhaar can lead to denial of welfare services, and at times exclusion

0.8% of people experienced exclusion due to Aadhaar-related reasons from a key welfare service (PDS, MGNREGS, social pensions) which they had earlier received. (Our survey also found that 3.3% of people experienced exclusion because of non-Aadhaar related problems from a key welfare service which they had earlier received).

1% of MGNREGS job card holders did not get work the last time they tried due to Aadhaar-related reasons (compared to 31% due to non-Aadhaar related reasons).

0.5% of social pension beneficiaries did not receive their pension the last time they expected it due to problems with Aadhaar (compared to 5.7% who did not receive it due to non-Aadhaar related reasons and many more who could not identify a reason).

1.5% of PDS users experienced a biometric authentication failure and did not receive ration in their last attempt. However, 3.2% of PDS users received their rations despite biometric authentication failure.

7. **Despite the Supreme Court ruling, many people find that Aadhaar is de facto mandatory for bank accounts, SIM cards, and school enrollment**

More than half of all people who produced Aadhaar to get a SIM card or bank account, said their provider accepted only Aadhaar, even after the Supreme Court ruling.

3.3% of people were denied bank accounts, and 0.8% of people were denied SIM cards due to Aadhaar-related problems.

0.5% of 6 to 14-year old's could not enroll in school due to Aadhaar-related reasons.

65% of people mistakenly believe that providing Aadhaar is mandatory by law for bank accounts, SIM cards, and school enrollment.

8. **Most people appreciate Aadhaar's universal acceptance; some have concerns**

72% of residents appreciate the convenience of Aadhaar as a universal ID, but almost **half** of these same people worry about the risks of linking it to too many services.

92% of people are very satisfied or somewhat satisfied with Aadhaar.

67% of people who have been excluded from a service due to problems with Aadhaar still say they are satisfied.

9. **The newer digital features of Aadhaar are yet to be embraced**

77% of people have never used a newer digital feature of Aadhaar (such as the mAadhaar app, QR code scanning, virtual Aadhaar number, or masked Aadhaar).

Only 39% have linked a correct mobile phone number to their Aadhaar.

10. Most people trust the Aadhaar system

90% of people trust that their data are safe in the Aadhaar system.

61% of welfare beneficiaries trust that Aadhaar prevents others from accessing their benefits.

8% worry about the misuse of their Aadhaar and 2% have experienced fraud that they believe was related to Aadhaar.

10 Other Empirical Evidence

There is also some rigorous empirical evidence that documents the welfare improvements caused by the use of biometric identity such as Aadhaar and we focus on three such areas.

● Public Distribution Programs

There is burgeoning empirical evidence on the effectiveness of public distribution programs linked to some form of identity. Banerjee, Duflo, Imbert, Mathew, and Pande (2020) find that e-governance led to reduction of leakage of public funds but not improvement of delivery of public services. In contrast, Muralidharan, Niehaus, and Sukhtankar (2016) find that biometric identification improved beneficiary experiences in collection and reduced corruption without incurring a substantial cost for the state. Bossuroy, Delavallade, and Pons (2019) find that biometric tracking is useful in delivering healthcare in challenging areas. Agarwal, Prasad, Sharma, and Tantri (2018) find that digital identity card-based workfare programs make these programs counter-cyclical, thus fulfilling one of the stated objectives of these poverty alleviation programs.

● Financial Inclusion

Empirical evidence for usage of banking services relies mostly on micro-evidence (Cole, Sampson, and Zia 2011; Dupas, Dean, Robinson, and Diego 2018; Dupas and Robinson 2013; Jack and Suri 2014; and Prina 2015). Aker, Boumniel, McClelland, and Tierney (2016) find that households

in Niger receiving mobile transfers had higher diet diversity and children consumed more meals per day which could partially be attributed to the time saved in travelling and waiting to receive their transfers. Giné, Goldberg, and Yang (2012) through a field experiment in Malawi find that fingerprinting borrowers when applying to loans led to substantially higher repayment rates for borrowers with the highest ex ante default risk but had no effect for the rest of the borrowers. Bruhn and Love (2014) and Burgess and Pande (2004) analyze the impact of access to banking on aggregate income and labor markets. Agarwal, Alok, Ghosh, Ghosh, Piskorski, and Seru (2017) study a large-scale experiment (related to Aadhaar) that led to 255 million new bank account opening in India. They find an increase in lending and default rate on new loans in regions with low ex-ante access to banking services. Balasubramanian, Chandra, Murlidharan, and Tantri (2019) document that “transaction failures do not deter the users from using the biometric platform” and that “convenience offered by technology seems to score over other conventional banking channels, despite high failure rate.”

- **Governance**

Digital identity can also help in improved governance. For example, Duflo, Hanna, and Ryan (2012) find that time-stamped photos and monetary incentives increased teacher attendance. There is also anecdotal evidence (Mookerji and Agarwal 2015) that Aadhaar resulted in government employees spending on average 20 min more daily at their work desk. Aadhaar enabled biometric attendance system has been implemented in the public healthcare system to address the problem of absent doctors (Yasmeen 2019). However, research on whether these initiatives have had an effect on performance is scant.

11 Other Identities

Many countries have closely watched the adoption and implementation of India’s ambitious identity system Aadhaar. There is a lot of interest in adopting a digital identity system to realize the benefits provided by digital architecture. At the same time, the ambitious nature of a centralized system and the privacy threats it engenders are making nations think carefully about data privacy and security issues (Rao and Nair (2019)). The Modular Open Source Identity Platform (MOSIP) developed in collaboration with many volunteers who were involved in the development of Aadhaar, anchored at the International Institute of Information Technology (IIIT), Bangalore,

integrates the lessons from Aadhaar and other identity systems around the world. An alternative to a centrally planned ID system is to let the private sector develop their own solutions and protocols that are not centralized (Allu, Deo, and Devalkar 2019). Many economists, including Nobel Laureate Paul Romer (India Today 2018), agree that establishing a robust, secure, and efficient identity protocol is vital (Carstens 2019). Exactly what form this will take will depend on each country's priorities and development goals (Economist 2018b).

Acknowledgements Bhagwan Chowdhry is professor of finance and Executive Director, Digital Identity Research Initiative (DIRI) funded by Omidyar Network at the Indian School of Business (ISB) and research professor at UCLA Anderson School. Amit Goyal is professor of finance at the Swiss Finance Institute of the University of Lausanne. Syed Anas Ahmed is research associate at DIRI at ISB. We thank Subhashish Badra and C.V. Madhukar of Omidyar Network, India and Sanjiv Das from Santa Clara University for carefully reading earlier drafts and providing thoughtful comments.

References

- Agarwal, Shweta, Subhashish Banerjee, and Subodh Sharma. 2017. Privacy and security of aadhaar: A computer science perspective. *Economic and Political Weekly* 52 (37): 93–102.
- Agarwal, Sumit, Shashwat Alok, Pulak Ghosh, Soumya Ghosh, Tomasz Piskorski, and Amit Seru. 2017. Banking the unbanked: What do 255 million new bank accounts reveal about financial access? Working paper available at SSRN: <https://ssrn.com/abstract=2906523>.
- Agarwal, Sumit, Shradhey Parijat Prasad, Nishka Sharma, and Prasanna L. Tantri. 2018. A friend indeed: Does the use of biometric digital identity make welfare programs counter cyclical? Working paper available at SSRN: <https://ssrn.com/abstract=3201746>.
- Aghion, Philippe, Patrick Bolton. 1997. A theory of trickle-down growth and development. *Review of Economic Studies* 64 (2): 151–172.
- Aiyar, Shankar. 2017. Aadhaar a biometric history of india's 12-digit revolution, *Westland*.
- Aker, Jenny C., Rachid Boumnijel, Amanda McClelland, and Niall Tierney. 2016. Payment mechanisms and anti-poverty programs: Evidence from a mobile money cash transfer experiment in Niger. *Economic Development and Cultural Change* 65 (1): 1–37.

- Allu, Rakesh, Sarang Deo, and Sripad Devalkar. 2019. Alternatives to aadhaar-based biometrics in the public distribution system. *Economic and Political Weekly* 54 (12): 30–37.
- Balasubramanian, Padmanabhan, S.V. Ravi Chandra, Aditya Murlidharan, and Prasanna L. Tantri. 2019. Fintech for the poor: Do technological failures deter financial inclusion? Working paper available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3840021.
- Banerjee, Abhijit, and Andrew F. Newman. 1993. Occupational choice and the process of development. *Journal of Political Economy* 101 (2): 274–298.
- Banerjee, Abhijit, Esther Duflo, Clement Imbert, Santhosh Mathew, and Rohini Pande. 2020. E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in India, forthcoming. *American Economic Journal: Applied Economics*.
- Besley, Timothy, and Torsten Persson. 2009. The origins of state capacity: Property rights, taxation, and politics. *American Economic Review* 99 (4): 1218–1244.
- Besley, Timothy, and Torsten Persson. 2010. State capacity, conflict, and development. *Econometrica* 78 (1): 1–34.
- Bossuroy, Thomas, Clara Delavallade, and Vincent Pons. 2019. Biometric tracking, healthcare provision, and data quality: Experimental evidence from tuberculosis control, NBER working Paper 26388.
- Bruhn, Miriam, and Inessa Love. 2014. The real impact of improved access to finance: Evidence from Mexico. *Journal of Finance* 69 (3): 1347–1376.
- Burgess, Robin, and Rohini Pande. 2004. Do rural banks matter? Evidence from the Indian social banking experiment. *American Economic Review* 95 (3): 780–795.
- Carstens, Agustin. 2019. Central banking and innovation: Partners in the quest for financial inclusion, seventeenth C.D. Deshmukh Memorial Lecture, April 25.
- Chaum, David, Debajyoti Das, Farid Javani, Aniket Kate, Anna Krasnova, Joeri De Ruiter, and Alan T. Sherman. 2017. cMix: Mixing with Minimal Real-Time Asymmetric Cryptographic Operations, In *Applied Cryptography and Network Security*, eds. Gollmann D., Miyaji A., Kikuchi H. ACNS 2017. Lecture Notes in Computer Science, vol 10355. Cham: Springer.
- Cole, Shawn, Thomas Sampson, and Bilal Zia. 2011. Prices or knowledge? What drives demand for financial services in emerging markets? *Journal of Finance* 66 (6): 1933–1967.
- Dalberg. 2019. State of aadhaar report 2019, Dalberg.
- Duflo, Esther, Rema Hanna, and Stephen P. Ryan. 2012. Incentives work: Getting teachers to come to school. *American Economic Review* 102 (4): 1241–1278.
- Dupas, Pascaline, and Robinson Jonathan. 2013. Why don't the poor save more? Evidence from a health savings experiment. *American Economic Review* 103 (4): 1138–1171.
- Dupas, Pascaline, Karlan Dean, Robinson Jonathan, and Ubfal Diego. 2018. Banking the unbanked? Evidence from three countries. *American Economic Journal: Applied Economics* 10 (2): 257–297.

- Economist. 2018a. A watertight store of Indians' personal data proves leaky, January 11.
- Economist. 2018b. Establishing identity is a vital, risky and changing business, December 18.
- Gelb, Alan and Julia Clark. 2013. Identification for development: The biometrics revolution. Working paper 315, Centre for Global Development.
- Giné, Xavier, Jessica Goldberg, and Dean Yang. 2012. Credit market consequences of improved personal identification: Field experimental evidence from Malawi. *American Economic Review* 102 (6): 2923–2954.
- Gunipati, Thulasiram. 2017. Know all about the backbone of aadhaar—Big data! *upGrad Blog*, November 14.
- India Today. 2018. When Nobel winner Paul Romer praised India's aadhaar scheme, October 8.
- Jack, William, and Tavneet Suri. 2014. risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review* 104 (1): 183–223.
- Khera, Reetika. 2011. The UID project and welfare schemes. *Economic and Political Weekly* 46 (9): 38–43.
- Khera, Reetika. 2018. These digital IDs have cost people their privacy— and their lives. *Washington Post*, August 9.
- Khera, Reetika. 2019. Aadhaar failures: A tragedy of errors. *EPW Engage*, April 5.
- Kremer, Michael. 1993. The O-ring theory of economic development. *Quarterly Journal of Economics* 108 (3): 551–575.
- Krusell, Per, and José-Victor Ríos-Rull. 1996. Vested interests in a positive theory of stagnation and growth. *Review of Economic Studies* 63 (2): 301–329.
- Leff, Nathaniel H. 1964. Economic development through bureaucratic corruption. *American Behavioral Scientist* 8 (3). 8–14.
- Mookerji, Nivedita, and Surabhi Agarwal. 2015. More than 30% govt staff clocked in before time on new year, *Business Standard*.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. Building state capacity: Evidence from biometric smartcards in India. *American Economic Review* 106 (10): 2895–2929.
- Niehaus, Paul, and Sandip Sukhtankar. 2013. Corruption dynamics: The golden goose effect. *American Economic Journal: Economic Policy* 5 (4): 230–269.
- Nilekani, Nandan. 2018. Data to the people. *Foreign Affairs*, September/October.
- OECD. 2018. Embracing Innovation in Government—Global Trends 2018, OECD.
- Pande, Varad. 2014. Direct benefits transfer: An idea whose time has come. *Livemint*, March 14.
- Parker, Ian. 2011. The I.D man. *New Yorker*, September 26.
- Prina, Silvia. 2015. Banking the poor via savings accounts: Evidence from a field experiment. *Journal of Development Economics* 115 (C): 16–31.
- Rangarajan, C. 2008. Report of the committee on financial inclusion. Ministry of Finance, Government of India.

- Ranjan, Mukesh. 2018. Social economist Jean Dreze blames aadhaar based authentication in PDS for seven starvation deaths in Jharkhand. *Indian Express*, June 22.
- Rao, Ursula, and Vijayanka Nair. 2019. Aadhaar: Governing with biometrics. *South Asia: Journal of South Asian Studies* 42 (3): 469–481.
- Sen, Amartya. 2000. Social exclusion: Concept, application, and scrutiny. In *Social development papers 1*. Manila: ADB.
- Singh, Pawan. 2019. Aadhaar and data privacy: Biometric identification and anxieties of recognition in India, forthcoming. *Information, Communication & Society*.
- Sriram, M.S. 2014. Identity for inclusion: Moving beyond aadhaar. *Economic and Political Weekly* 49 (28): 148–154.
- UIDAI. 2017. Introduction to aadhaar, Unique Identification Authority of India.
- Varma, Pramod K. 2010. Aadhaar Scalability and Data Management Challenges, UIDAI, https://www.cse.iitb.ac.in/~comad/2010/pdf/Industry%20Sessions/UID_Pramod_Varma.pdf.
- Wimmer, Kurt, and Gabe Maldoff. 2019. India proposes updated personal data protection bill. *Inside Privacy*, December 12.
- World Bank. 2016. ID for Development (ID4D) Report, Washington DC, World Bank.
- Yang, Dean. 2008. Can enforcement backfire? Crime displacement in the context of customs reform in the Philippines. *Review of Economics and Statistics* 90 (1): 1–14.
- Yasmeen, Afshan. 2019. Aadhaar-enabled biometric attendance system for health department staff. *Hindu*, October 13.



Big Techs in Finance: On the New Nexus Between Data Privacy and Competition

Frederic Boissay, Torsten Ehlers, Leonardo Gambacorta,
and Hyun Song Shin

1 Introduction

Large technology companies such as Alibaba, Amazon, Facebook, Google and Tencent have begun to enter financial services. Their entry builds on their established digital platforms in e-commerce, search and social media, and holds the prospect of efficiency gains and greater financial inclusion (BIS, 2019). The business model of these “big techs” rests on enabling direct interactions among a large number of users. An essential by-product of their business is their large stock of user data, which are used as an input for a range of services that exploit natural network effects, generating further user activity. Increased user activity then completes the circle, as it generates yet more data. The self-reinforcing loop between Data, Network externalities and Activities, is the DNA of big techs.

F. Boissay · T. Ehlers · L. Gambacorta · H. S. Shin (✉)
Bank for International Settlements, Basel, Switzerland
e-mail: hyunsong.Shin@bis.org

F. Boissay
e-mail: frederic.boissay@bis.org

T. Ehlers
e-mail: torsten.ehlers@bis.org

L. Gambacorta
e-mail: leonardo.gambacorta@bis.org

Personal data lie at the heart of this new, digital economy. Building on the advantages of the self-reinforcing nature of the data-network-activities loop, some big techs have ventured into financial services, including payments, money management, insurance and lending. As yet, financial services are only a small part of big techs' business globally. But given their size and customer reach, big techs' entry into finance has the potential to spark rapid change in the financial industry. The entry of big techs promises many potential benefits. With their low-cost structures, big techs can easily scale up their businesses to provide basic financial services, especially where a large part of the population remains unbanked. Using big data and analysis of the network structure of user connections in their platforms, big techs can better assess the riskiness of borrowers, reducing the need for collateral to assure repayment. As such, big techs promise to enhance the efficiency of financial services provision, promote financial inclusion and allow associated gains in economic activity.

The benefits notwithstanding, widespread use of data in combination with the new technologies and applications can have adverse side effects. Big techs have the potential to become dominant through the advantages afforded by the data-network-activities loop, raising competition and data privacy issues. Complex questions arise about how best to organise access to personal data, not only to limit side effects, but also to protect people's privacy to the degrees desired. How to define and regulate the use of data has become an important policy issue for authorities at the domestic and international level.

This chapter starts by explaining big techs' business model, their life cycle and the characteristics of the data-network externalities-activities feedback loop. A second section analyses how big techs' DNA could produce potential benefits in the provision of financial services, both increasing financial inclusion and reducing asymmetric information problems in the supply of credit. A third section evaluates the potential costs caused by the entry of big techs into finance, including the new risks of price discrimination, abuse of market power, anti-competitive behaviour and limits to data privacy. A last section lays out the complex public policy trade-off between the objectives of efficiency and privacy, and discusses the policy options, as well as the case for policy coordination at the domestic and international level.

2 Big Techs' Business Model

Big techs' business models can be best described as online platforms that allow different types of users (e.g. buyers and sellers) to interact. This

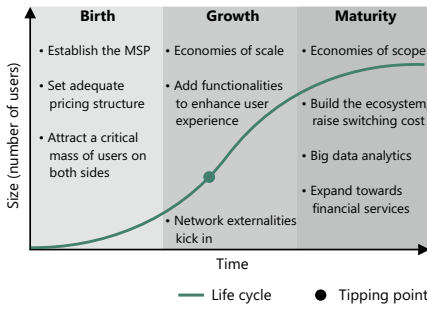
creates network externalities: the more users interact, the more attractive the platform. Other industries (e.g. telecommunication networks, credit card payments networks, etc.) feature network externalities. But big techs' online-focused business models allow them to reach dominant market positions at unprecedented speed. Further, the systematic accumulation of user data and new ways of analysing it (e.g. artificial intelligence such as machine learning solutions) allows them to exploit these network externalities in a very effective way.

2.1 Big Techs as Multi-Sided Platforms and Their Life Cycle

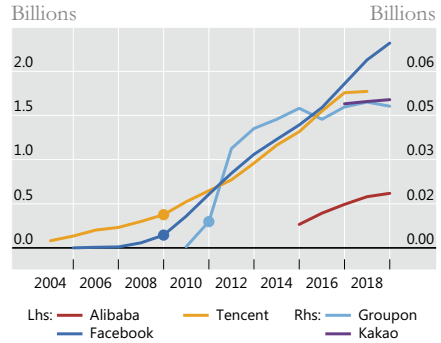
Big techs initially create value as online “multi-sided platform s” (MSPs), by enabling and catalysing direct interactions between two or more groups of users (e.g. buyers and sellers). The three main types of online platform are social networks, e-commerce platforms and search engines.

In contrast to traditional bilateral exchanges, users on each side transact with each other through the platform—not necessarily with the platform itself. Social platforms, for example, allow people to connect to each other, so that each member benefits from a larger community. E-commerce websites enable their users to buy and sell a wide variety of goods and services worldwide. The larger number of sellers reduces buyers' search costs, and a larger number of buyers expands sellers' business opportunities. A typical feature of MSPs is the presence of network externalities: the very fact that users participate on one side of the platform (e.g. buyers) increases users' benefits on the other side (e.g. sellers). One challenge is to attract users onto both sides at the same time—a “chicken and egg” problem. Successful platforms solve this problem by using specific price structures, which essentially consist in charging a lower fee to the side that creates the most network externalities—and letting the side that benefits the most from the network subsidise the other (see, e.g., Jullien 2004).

Big techs have so far followed a rather traditional corporate life cycle with three phases: birth, growth and maturity (Graph 1, left-hand panel). What makes them unique is the coincidence of several factors (i.e. the collection of personal data on a large scale, network effects and a large number of activities) and the rapidity with which they reach maturity. Petralia et al. (2019), for example, report that social networks such as Facebook or Tencent's WeChat took less than five years to reach 50 million users (see also Graph 1, right-hand panel). In terms of user numbers, these firms are much larger, and have grown much faster, than any financial firm. In particular, Nguyen Trieu

From big techs' birth to maturity¹

Big techs' monthly active users



Graph 1 Big techs' life cycle: Theory and practice. ¹The firm's life cycle described in the left-hand panel borrows from the synthesis of the literature by Miller and Friesen (1984). Given that big techs are still new and rising firms, we purposely ignore the usual "decline" phase. MSP = multi-sided platform (Source BIS [2019]; Miller and Friesen [1984]; S&P Capital IQ; Authors' calculations)

(2017) argues that big techs have scaled up between 10 and 100 times faster than traditional financial institutions.

Even after an MSP has attracted a sufficient mass of users on both sides, the emphasis remains on further increasing the number of users, with the aim of reaching a tipping point where adoption rates accelerate and network effects kick in. Beyond this point, growth can be very fast. More buyers bring more sellers—and vice versa—so that the MSP enjoys increasing returns to scale. The average cost of serving a user falls as the total number of users rises. And users are willing to pay more for access to a bigger network. As a result, the platform's margins improve.

2.2 The DNA of Big Techs

Data analytics, network externalities and interwoven activities ("DNA") comprise the key features of big techs' business models. These three elements reinforce each other. Network externalities generate additional users and added value for users. They allow the big tech to generate more data, the key input into data analytics. The analysis of large troves of data enhances existing services and attracts further users. More users, in turn, provide the critical mass of customers, so that a wider range of activities can be offered, which yields even more data. Accordingly, network externalities are stronger on platforms that offer a broader range of services, thus representing an essential element of big techs' life cycle.

Financial services both benefit from and fuel the DNA feedback loop. Offering financial services can complement and reinforce big techs' commercial activities. The typical example is payment services, which facilitate secure transactions on e-commerce platforms, or make it possible to send money to other users on social media platforms. Payment transactions also generate data detailing the network of links between fund senders and recipients. These data can be used both to enhance existing (e.g. targeted advertising) and other financial services, such as credit scoring.

The source and type of data and the related DNA synergies vary across big tech platforms. Those with a dominant presence in e-commerce collect data from vendors, such as sales and profits, combining financial and consumer habit information. Big techs with a focus on social media collect data on individuals and their preferences, as well as their network of connections. Big techs with search engines do not observe connections directly, but typically have a broad base of users and can infer their preferences from their online searches.

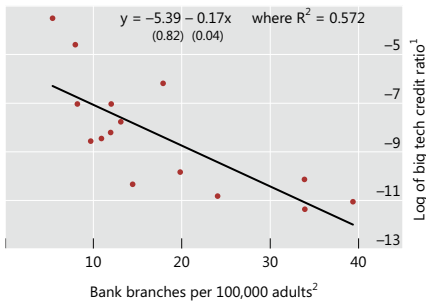
The type of synergy varies with the nature of the data collected. Data from e-commerce platforms can be a valuable input into credit scoring models, especially for SME and consumer loans. Big techs with a large user base in social media or internet search can use the information on users' preferences to market, distribute and price third-party financial services (e.g. insurance).

3 Potential Benefits from Big Techs

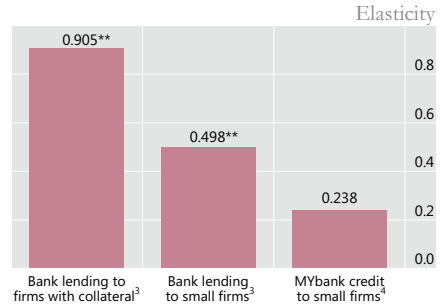
Building on the advantages of the reinforcing nature of the data-network-activities loop, some big techs have ventured into financial services, including payments, money management, insurance and lending. As yet, financial services are only a small part of their business globally (around 11% of their total revenues). But given their size and customer reach, big techs' entry into finance has the potential to spark rapid change in the industry. It offers many potential benefits. Thanks to their low-cost structures, the businesses of big techs can easily be scaled up to provide basic financial services, especially in places where a large part of the population remains unbanked.

Payments were the first financial service to be offered by big techs, mainly to help overcome the lack of trust between buyers and sellers on e-commerce platforms. Buyers want goods to be delivered, but sellers are only willing to deliver after being assured of payment. Payment services, such as those provided by Alipay (owned by Alibaba) or PayPal (owned by eBay), allow guaranteed settlement at delivery and/or reclaims by buyers and are fully

Big tech credit and banking sector development



Elasticity of credit with respect to house prices in China



Graph 2 Big tech credit, asset prices and bank development. Robust standard error in parentheses. **indicates significance at the 5% level. ¹The ratio is calculated for 2017 and is defined as big tech credit divided by total credit to the private non-financial sector (including total fintech credit). ²Average over the period 2010–2015. ³Period of estimation: 2005–2013. ⁴Period of estimation: 2011–2017

(Source BIS [2019]; Frost et al. [2019]; Gambacorta et al. [2019]; World Bank; Cambridge Centre for Alternative Finance and research partners; company reports; Authors' calculations)

integrated into e-commerce platforms. In some regions with less developed retail payment systems, new payment services have emerged through mobile network operators (e.g. M-Pesa in several African countries). Over time, big techs' payments services have become more widely used as an alternative to other electronic payment means such as credit and debit cards.

Benefits are also evident in lending provision. Using big data and by analysing the network structure of user connections in their platforms, big techs can assess the riskiness of borrowers, reducing the need for collateral to assure repayment. As such, big techs can promote financial inclusion.

3.1 Financial Inclusion

Financial inclusion allows individuals and businesses to have access to financial products and services in a responsible and sustainable way. In particular, the provision of credit or saving products to unbanked individuals allows them to use resources more optimally over time. Meanwhile, insurance products can serve as a cushion against shocks such as bad harvests, illness or the death of the main wage earner.

The greater use of data can foster greater convenience, more tailored and personalised products and greater financial inclusion. Big techs may have a competitive advantage over banks and serve firms and households that would otherwise remain unbanked (Graph 2, left-hand panel).¹ This is because they

can tap different but relevant information from their digital platforms.² For example, Ant Financial and Mercado Libre claim that their credit quality assessment and lending decisions typically involve more than 1,000 data series per loan applicant.

There is evidence that the advent of fintech and big tech lenders and their use of alternative data have been a boon for borrowers who are unserved or underserved by banks. In China, for example, the major platforms have provided access to credit for hundreds of millions of new personal and business borrowers. In many emerging markets, the use of data on transactions, payment of utility bills, or platform reviews is driving greater access to and personalisation of financial services. These benefits exist even in countries with advanced systems of financial services provision. For example, in the United States, better use of personal transaction data promises to help the 45–60 million “thin credit file” Americans, i.e. those who have inadequate credit history, to obtain loans. In other words, the efficient use of data allows big techs to overcome some of the traditional information problems typically encountered by traditional intermediaries with respect to more opaque borrowers.

3.2 Reduction of Financial Frictions in Lending

The sheer amount of data collected by big techs and their intelligent use have the potential to reduce financial frictions, in particular, borrower screening, monitoring and collateral requirements. Potential borrowers who cannot be served by regular banks due to prohibitive administrative costs could potentially obtain credit on the basis of credit ratings or scores built on a broader set of data processed in novel ways.

3.3 Borrower Information and Screening

Information costs on credit markets may sometimes be so prohibitive that banks refrain from serving borrowers—or do so only at very high spreads. Big techs’ processing of large quantities of information (big data) using advanced analytical methods such as machine learning and network analysis (artificial intelligence) can reduce such costs. Big data relevant for financial services obtained directly from big tech platforms include (i) transactions

¹ More generally, big techs’ market penetration rate tends to be higher in areas where banks are absent or their branch networks sparser.

² See Hau et al. (2018) and Huang et al. (2018) for the case of China.

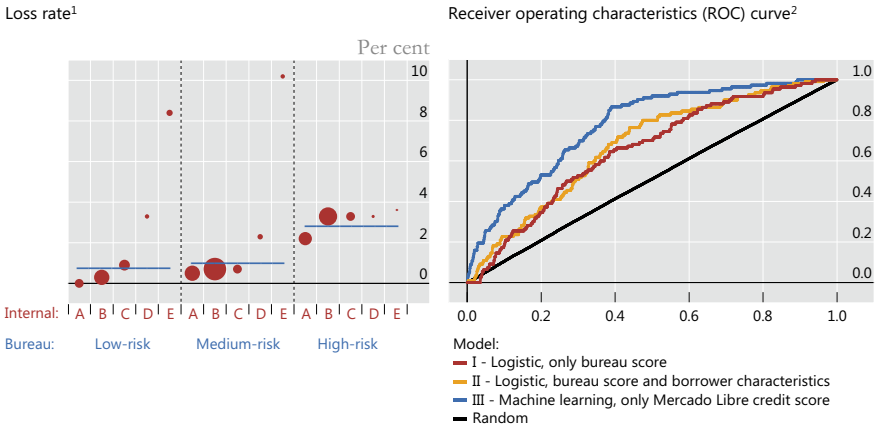
(sales volumes and average selling prices); (ii) reputation-related information (claim ratio, handling time, reviews and complaints); and (iii) industry-specific characteristics (sales seasonality, demand trend and macroeconomic sensitivity). This can be also enriched by using non-traditional data obtained via social media and other channels.

Frost et al. (2019) suggest that, when applied to small vendors, big techs' credit scoring outperforms models based on credit bureau ratings and traditional borrower characteristics. The predictive power of the big techs' scoring systems arises in large part from exploiting the network structure. For instance, MYbank (Ant Financial group) uses network analysis of transactions to evaluate whether an entrepreneur separates personal funds from business funds, which is one of the basic principles of good business conduct.

Frost et al. (2019) also present empirical evidence that the use of more granular data with machine learning can help to improve the predictive power of prepayment prospects, especially for small merchants who are typically not served by banks. In the case of Mercado Libre, internal ratings are more granular (A to E) than those of the credit bureaus in Argentina (low risk to high risk). Banks rely on information from credit bureaus but augment it with other borrower characteristics and soft information (Graph 3, left-hand panel). However, as most of Mercado Libre's clients are unbanked, the analysis below is more specific to cases in which traditional soft information collected by banks is not available.

For a given bureau rating (e.g. low risk), the expected loss rate is strictly monotonic with the internal rating (i.e. the patterns of the dots show that the internal rating increases with expected loss). Conversely, for a given internal rating (e.g. C, D or E), the loss rate is not strictly monotonic with the credit bureau risk. For example, the dot associated with internal rating D in the low-risk bureau category indicates a higher risk than the internal rating D in the medium-risk bureau category. Moreover, the internal rating has a broader range, covering losses from 0.0% to 10.2%; the bureau rating ranges from 0.7% to 2.8%.

Most importantly, by using the internal scoring model, Mercado Libre is able to provide credit to the profiles assessed as high risk by the bureau. The size of the dots in the left-hand panel of Graph 3 is proportional to the share of the firms in the rating distribution; a substantial number of clients are in the credit bureau high-risk category. Because banks use a mix of credit bureau information, hard information from financial statements and soft information from loan officers, this segment may have much less access to traditional banking services. With its more granular scoring model, Mercado Libre offers 30% of its credit to this category.



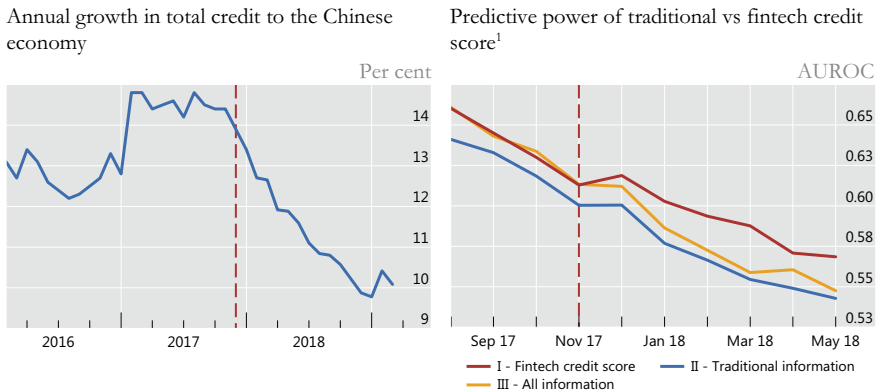
Graph 3 Credit assessment and big data analytics. ¹The loss rate is the volume of loans more than 30 days past due relative to the origination volume. To date, Mercado Libre's internal rating system has proved better able to predict such losses. It segments loan originations into five different risk groups as compared with the three clusters identified by the bank bureau. The size of the dots is proportional to the share of the firms in the rating distribution. ²True positive rates versus false positive rates for borrowers at different thresholds for a logistic model with only the credit bureau score (I), a logistic model with the bureau score and borrowers' characteristics (II), and a machine learning model with the Mercado Libre credit score (III). A random model is included for comparison purposes. The ROC curve shows that the machine learning model has superior predictive power to both the credit bureau score only and the credit bureau score with borrower characteristics (Source BIS [2019]; Frost et al. [2019])

Further, the internal rating system based on machine learning techniques and data obtained from the e-commerce platform can outperform simple models based on bureau score and borrower characteristics in predicting defaults (Graph 3, right-hand panel). Performance is measured in this case on the y-axis directly by the area under the receiver operating characteristics (AUROC) curve. The higher the AUROC, the higher the predictive power of the model. The right-hand panel of Graph 3 reports the respective performance of the three models using this measure. The AUROC ranges from 50% (purely random prediction) to 100% (perfect prediction). From Graph 3 it is evident that the predictive power rises the most for the model that uses a machine learning technique applied to the data from the e-commerce platform.

While encouraging, these results prompt the questions (i) whether the performance of machine learning or fintech-based credit scoring models is superior to bank models that also use soft information, and (ii) if their performance can be sustained over full business and financial cycles.

To answer these questions, Gambacorta et al. (2019) use a unique dataset on loan transactions from a leading Chinese fintech company to compare the predictive power of credit scoring obtained with machine learning and big data with more traditional linear models in the case of a shock. The predictive performance of the models in terms of credit losses and defaults is analysed both in “normal times” and following an (exogenous) change in regulatory policy on shadow banking in November 2017 (dashed line in the two panels of Graph 4). The new rules led traditional and new financial institutions to tighten their lending requirements, causing a significant drop in the growth of total credit (the left-hand panel of Graph 4) and an increase in the number of defaults. As in Frost et al. (2019), performance is measured directly by the area under AUROC curve (y-axis).

In normal times the fintech credit score outperforms the linear model with traditional information, but performs in line with the linear model that uses



Graph 4 Fintech credit scores show greater predictive power after a change in regulation. The vertical dashed line indicates when the People’s Bank of China (PBoC) issued specific draft guidelines to tighten regulations on shadow banking. In particular, from 17 November 2017, financial institutions have not been allowed to use asset management products to invest in commercial banks’ credit assets or provide “funding services” for other institutions (such as fintech companies) to bypass regulations. The new rule has had a huge impact on fintech companies’ funding sources. The PBoC set also a limit on the interest rates charged by P2P lending companies. All annualised interest rates, which include the upfront fees charged for loans, were capped at 36%. The effects of these new rules were also reinforced by the strict measures concerning online micro-lending that were imposed on December 1, 2017 by China’s Internet Financial Risk Special Rectification Work Leadership Team Office.

¹The vertical axis reports the Area Under the ROC curve (AUROC) for every model. The AUROC is a widely used metric for judging the discriminatory power of credit scores. The AUROC ranges from 50% (purely random prediction) to 100% (perfect prediction)

(Source Gambacorta et al. [2019])

all information, i.e. both traditional and non-traditional data (right-hand panel of Graph 4). However, the fintech credit score outperforms the other models in the aftermath of the modification in credit conditions that followed the change in regulation. One possible interpretation of this finding is that credit scoring models based on machine learning better capture the non-linear relationship between variables in a period of stress.

3.4 Monitoring and Collateral

The cost of enforcing loan repayments is an important component of total financial intermediation cost. To reduce enforcement problems banks usually require borrowers to pledge as collateral tangible assets, such as real estate, with the aim of increasing recovery rates in case of default. Another precaution is monitoring. Banks spend time and resources monitoring their clients' projects to limit the risk that borrowers implement them differently from what was initially agreed. Through the monitoring process, firms and financial intermediaries also develop long-term relationships and build mutual trust, which makes default a less attractive option for borrowers.

Big techs can address these issues differently. When a borrower is closely integrated into an e-commerce platform, for example, it may be relatively easy for a big tech to deduct the (monthly) payments on a credit line from any of the borrower's revenues that flow through its payment account. In contrast, banks may not be in the position to do so as the borrower may have accounts with other banks. Given network effects and high switching costs, big techs could also enforce loan repayments by the simple threat of a downgrade or an exclusion from their ecosystem if borrowers default on their payments. Anecdotal evidence from Argentina and China suggests that the combination of massive amounts of data and network effects may allow big techs to mitigate the information and incentive problems that are traditionally addressed through the posting of collateral. This could explain why, unlike the supply of corporate loans from banks, that of big techs does not closely correlate with asset prices (Graph 2, right-hand panel).

4 Potential Costs of Big Techs' Use of Personal Data

Big techs' entry into finance brings efficiency gains and lowers barriers to the provision of financial services. But the very features that bring benefits also have the potential to generate new risks and costs associated with

market power. Once a captive ecosystem is established, platforms can exploit their market power and network externalities to increase user switching costs, exclude potential competitors and consolidate their position by raising barriers to entry.

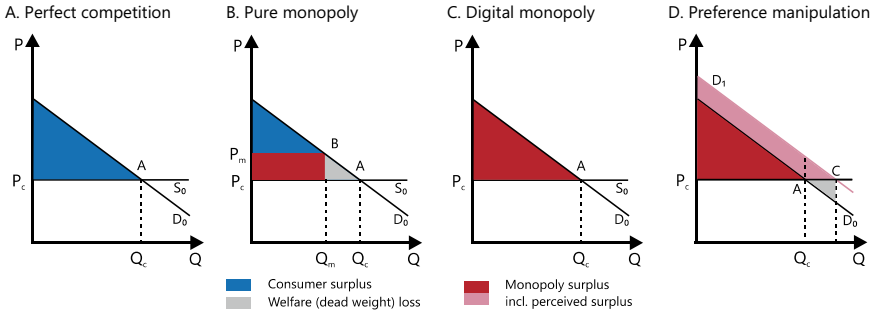
4.1 Monopolistic Use of Data and Price Discrimination

A major risk is the monopolistic use of data. One special aspect of data as an input of production is non-rivalry: data can be used many times over, and by any number of firms simultaneously, without being depleted. Thanks to non-rivalry, data generate increasing returns in both scale and scope (Farboodi et al. 2019). Any single additional piece of data (e.g. a credit score) has additional value when it is combined with an existing large stock of data or across sources. For this reason, data are more valuable to big tech firms, giving rise to feedback loops and so-called digital monopolies.

Data can be used for price discrimination. Once their dominant position in data is established, digital monopolies may use the data not only to assess a potential borrower's creditworthiness or riskiness, but also to identify the highest rate a given borrower would be willing to pay for a loan or the highest premium a client would pay for insurance (i.e. their individual "reservation price"). Fuelled by big data, algorithmic price discrimination is able to parse the population of potential customers into finer and finer subcategories—each matched with a different price. In some cases, sellers are even able to set personalised pricing, marching down the demand curve and setting a different price for each consumer.³

A significant source of profits for big techs comes from this extraction of the consumer's surplus. Bar-Gill (2019) illustrates this mechanism by comparing the equilibrium outcomes under perfect competition (Graph 5, panel A), pure monopoly (i.e. without the use of big data; panel B), and digital monopoly (panel C). Under perfect competition, financial services are priced at their marginal cost (P_c), and the entire surplus accrues to consumers (blue area). The pure monopoly also sets one single price but, as this price is typically higher than the marginal cost, the supply of services is reduced. The overall surplus shrinks (by an amount that corresponds to the grey area), and the monopoly corners some of it (red area). In this case, consumers are worse off. Panel C presents the case of a *digital* monopoly. Using big data and

³ There is empirical evidence of price discrimination based on information collected online about consumers. The price difference for identical products may vary up to 30%, depending on the location and the characteristics (for instance, browser configurations) of different online visitors. See Mikians et al. (2013) and Bar-Gill (2019).



Graph 5 From market competition to market manipulation (Source Adapted from Bar-Gill [2019])

sophisticated algorithms, a monopolistic big tech identifies each consumer’s reservation price, and sets a personalised price just below it. Price discrimination allows the big tech to increase the quantity sold, back up to the competitive outcome (i.e. from point B to point A). The deadweight loss is eliminated, which increases the overall social welfare back to the perfect competition level. However, the entire surplus now goes to the big tech. In that case, efficiency is restored at the cost of the redistribution of the surplus in favour of the digital monopoly. But, ultimately, consumers are worse off than under a pure monopoly.

4.2 Algorithmic Biases and Abuse of Market Power

The algorithms used to process data may also develop biases, leading to unethical discrimination (e.g. based on race, religion, etc.) and greater inequality (O’Neil 2016). For instance, one recent study of the US mortgage market found that black and Hispanic borrowers were less likely to benefit from lower interest rates from machine learning-based credit scoring models than non-Hispanic white and Asian borrowers (Fuster et al. 2019).

Even more worrying is the potential for intentional manipulation. Evidence suggests that big tech firms can exploit consumers’ behavioural biases. For instance, one experiment based on about 670,000 unaware Facebook users found that people’s emotional state can be projected onto others through contagion. This ability to make people experience the same emotions without being aware of the cause clearly raises economic, not to mention ethical, concerns (Kamer et al. 2014).⁴

⁴The notion that firms may actively change preferences and create wants, e.g. through advertising and salesmanship, is already present in Galbraith (1958). But the scope for such actions may be

When consumer preferences can be manipulated, the loss of surplus for consumers can be large. Panel D in Graph 5 illustrates this point. It represents what happens when a digital monopoly persuades its consumers to overestimate the benefit from a service or product. Graphically, the consumer demand curve shifts eastward (from D_0 to D_1). The overestimation causes some consumers to purchase the product, even though its actual value to them is lower than the price. As the additional surplus is only *perceived* (light red area), these consumers, who purchase the product only because of their misperception of the benefit, suffer an even greater welfare loss (light red areas) than under price discrimination. This outcome too is inefficient (grey area).

4.3 Anti-Competitive Behaviour

Another potential market failure could arise from big techs' *control over key digital platforms*. Once a captive ecosystem is established, potential competitors face steep costs and high risks in setting up rival platforms.

On the one hand, the fixed cost of setting up a new network of users, for instance, can be prohibitive.⁵ This could allow big techs to engage in traditional anti-competitive practices such as “tying-in sales”, cross-subsidising activities, and purchasing competitors.⁶ On the other hand, big techs' search, mobile network, social network or e-commerce platforms have become essential facilities for an ever wider range of business activities.

Platforms now often serve as essential selling infrastructures for financial services providers, while at the same time big techs compete with these same providers. When a network operator owns a smartphone-based payment system, for example, it can undermine competitors' access to its own digital platform by charging competitors (i.e. banks or rival big techs) high fees to connect with its (payment) system. Similarly, the owner of a search engine may redirect users away from competitors towards their own brands. With

greater in the case of big techs, due to their command over much richer customer information and their integration into their customers' everyday life.

⁵ To date there is no evidence of big techs hindering their competitors' provision of financial services on their platforms. But examples of anticompetitive practices can be found in other sectors of activity, such as advertising. For example, in March 2019 the European Commission fined a big tech for imposing a number of restrictive clauses in contracts with third-party websites which prevented its rivals from placing their search adverts on those websites.

⁶ For example, Facebook acquired Instagram – a photo app – in 2012, Onavo – a data-security app that tracks users' smartphone activity – in 2013, WhatsApp – a messaging service – 2014; and Tbh – a social-polling app – in 2017. When Snapchat rebuffed its purchase offer in 2013, Facebook responded by cloning the app's most successful features.

a captive user base, dominant platforms can raise the price of their services, and extract a larger share of customer surplus.

4.4 Privacy

When information is gathered without the informed consent of the consumer, it often infringes on personal privacy. Popular health websites have been found to share people's sensitive data (e.g. medical symptoms, diagnoses, drug names, etc.) with dozens of companies around the world, including big tech firms such as Google, Amazon and Facebook (Financial Times 2019a). These risks are still greater when firms underinvest in data security, leading to data breaches (Carrière-Swallow and Haksar 2019).

Furthermore, armed with this knowledge, companies may be in a better position to sell specific treatments, services or financial products that may not be in the users' interests. Consumer welfare will not necessarily benefit from the collection and use of personal data by profit-oriented firms, if these are left to their own devices.

5 The New Policy Trade-off Between Efficiency and Privacy

The benefits and costs of the use of personal data in finance raises important policy questions. These go beyond the traditional ones of financial stability and competition, extending also to a new trade-off between data efficiency and privacy.⁷

5.1 Challenges for Public Policy

A first challenge is related to assigning the control and ownership of personal data. Control and ownership of data are rarely clearly defined. In some cases, users volunteer to provide their data for free (e.g. posts on social networks, online registrations). In other cases, companies infer personal information indirectly, e.g. through users' social networks or internet searches. Users may also unintentionally and inadvertently surrender data, e.g. through their digital footprints or geo-localisation. Ownership and control of data is also difficult to re-assign, e.g. to users, due to the wide variety of data and the ways data are

⁷ See BIS (2019) and Petralia et al. (2019).

gathered, as well as their intangibility and non-rivalry. In many countries, the default outcome is that financial institutions or big techs acquire customer data at a very low cost and keep de facto control.

A second challenge is addressing the heterogeneity in personal data. The efficiency gain from personal data-sharing crucially depends on their type. Some data are purely private or only meant to be shared with a restricted number of users—e.g. medical records. At the other extreme are data that people may want to share freely, and which can be shared without causing any harm. In between, there may be data that can be lent out (temporarily) and combined with other data, e.g. for credit assessments or insurance pricing. Some data are not valuable to users (e.g. browsing histories), but may be valuable to private sector companies, as they may help both general and customer-specific services to be better targeted. Users may want to sell such data to the highest bidder. But an efficient and complete market for personal data has not yet emerged. To establish such a market, one must first determine who controls the data.

A third challenge relates to the value of privacy, and whether privacy should be traded off against other goals in the first place. In assessing the cost of the widespread sharing and use of data, one needs to consider how much people value their privacy.⁸ Some argue that data privacy has the attributes of a fundamental right, which cannot be traded off against economic benefits.⁹

However, the evidence suggests that cultural preferences towards data privacy differ across jurisdictions, and even between different social segments (Chen et al. 2020). For example, in one recent survey, respondents were asked if they would be open to their bank securely sharing their data with other organisations in exchange for better offers on financial services.¹⁰ In India, 65% of respondents said yes. In the Netherlands, this was only 13%.

Overall, it appears that willingness to share data correlates with per capita income, declining as incomes increase (see Graph 6, left-hand panel). Within jurisdictions, there are large differences by age and gender. For instance, 38% of 25- to 34-year-olds globally were willing to share their data, but only 16% of those over 65 were. At the global level, 34% of men were willing to share data, but only 27% of women, with this gap even larger in some countries (see Graph 6, right-hand panel).

Another important aspect in valuing privacy is with whom people are willing to share their private data. A recent survey on the anonymous

⁸ For a further discussion, see Acquisti et al. (2016).

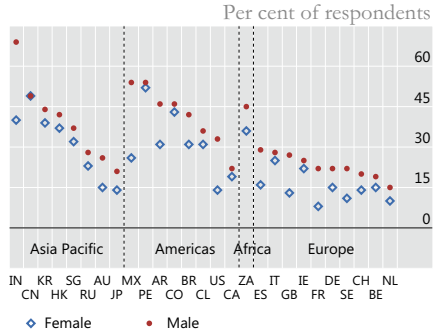
⁹ For a discussion of data rights in Europe, and of the grounding of such rights in e.g. the EU Charter of Fundamental Rights, see BEUC (2019).

¹⁰ See EY (2019). The survey covered 27,000 consumers in 27 markets.

Willingness to share data falls with GDP per capita



Willingness to share data is higher for men than women¹



Graph 6 Willingness of users to share data for better offers in financial services varies. ¹The question in the survey reads “I would be comfortable with my main bank securely sharing my financial data with other organisations if it meant that I received better offers from other financial intermediaries”; for Belgium the figure is calculated over Belgium and Luxembourg (Source EY [2019]; IMF, *World Economic Outlook*; Chen et al. [2020])

sharing of genetic and medical information points, again, to wide cross-country differences. For example, more than 65% of the Chinese and Indian respondents indicate that they would share their health data with governmental authorities. In contrast, the proportion is less than 20% in European countries.¹¹ Such varied preferences make it difficult to reach universally acceptable solutions when designing controls and pricing data.

5.2 Policy Options

There are several, potentially complementary, approaches to address the efficiency-privacy trade-off raised by the widespread use of personal data. One approach consists of restricting the processing of user data. For example, recent data protection laws (e.g. in the European Union, Brazil, California, Japan, Singapore) have clarified data collection and use to protect personally identifiable information. The challenge with these laws is how to balance differences in privacy concerns and use of data. For instance, India has instituted storage rules for payment system data motivated by privacy, access by regulators to necessary information and obligations to the judiciary in

¹¹ The Welcome Sander Institute, “Your DNA, Your Say” global online survey is an ongoing survey that gathers public attitudes towards genomic data sharing for over 37,000 individuals in 22 countries.

cases of legal disputes.¹² Some other jurisdictions have taken measures with a wider ambit that may restrict data flows across borders.¹³ The rationale for such measures relates to law enforcement and monitoring and supervision purposes. Yet, such frictions in the use of data could lead to cost inefficiencies, limiting their potential benefits.¹⁴

A second approach is to give consumers greater control over their personal data. This could allow customers to grant competitive firms access to relevant information, and can thus foster competition. Recent open banking initiatives (e.g. in the European Union, Australia, and Mexico) are examples of concrete policy actions in this direction. These policies require financial firms to make their customers' transactions (or equivalent) data portable, i.e. directly transferable to third parties or competitors, typically through open application programming interfaces (APIs). Open banking rules selectively restrict the range of data that can be transmitted (e.g. financial transaction data), as well as the type of institutions among which such data can be shared (e.g. accredited deposit-taking institutions). In this sense, they help resolve inefficiencies through the allocation of property rights and the creation of a competitive market for data—the decentralised or “Coasian” solution.¹⁵ As access is asymmetrical, though, open banking rules do not fully level the playing field between big tech firms and incumbent service providers.

A third approach is a set of public infrastructures on which a layer of services can be built. This includes important foundations for digital services such as digital identity, like Aadhaar in India and MyInfo in Singapore, and the development of data management protocols. Once these infrastructures are in place, payments, digital government services and a host of other solutions become possible. For example, the launch of India's Unified Payment Interface (UPI) facilitated entry by new firms and spurred competition, which drastically reduced prices for consumers.¹⁶ By making consumers data-rich and giving them greater control over their data, important benefits for users can be reaped (Nilekani 2018).

¹² India does not restrict the transfer of payment system data overseas if one of or more of the counterparties is a non-resident, or for processing purposes even when both counterparties are residents, but in the latter case local data storage is mandatory as the data belong to Indian citizens.

¹³ One example is China (see Cyberspace Administration of China 2019).

¹⁴ According to Aaronson (2019), 58% of the countries in the world have now adopted or are adopting data protection laws. Many such laws contain provisions affecting cross-border data flows. It is still too early to assess whether such laws are effective in addressing risks (see also Mitchell and Mishra 2019).

¹⁵ This is named after Ronald Coase; see Coase (1960).

¹⁶ See Financial Times (2019b).

5.3 Domestic and International Policy Coordination

How to define and regulate the use of data are issues that need to be coordinated at both the domestic and international level.

At the domestic level, central banks and financial regulators may need to upgrade their understanding of personal data issues. And they need to coordinate with competition and data protection authorities. Yet the mandates and practices of these bodies may not always be compatible. For example, the specificity of the financial sector may not accord with the general standards that competition and data privacy laws often apply to a wide range of industries. Moreover, financial regulation is often based on international standards, while data protection and competition policy are mostly national—to the extent that countries may have a unique competition or data protection authority.¹⁷

At the international level, regulations on the use of personal data diverge widely. In the European Union, the General Data Protection Regulation (GDPR) assigns data rights to individuals. In India, the India Stack generates large volumes of new data, and users have control over them, but data privacy regulations have to yet be adopted. In China and several other countries, data localisation rules prevent data from being shared across borders. In the United States, a patchwork of sector-specific legislation means that, in practice, companies have relatively free access to data, and some companies, most famously Apple, have resisted calls to share data with public authorities (Apple 2016). Meanwhile, only a few countries have a national data or artificial intelligence strategy.¹⁸ As the digital economy expands across borders, there is a need for international cooperation on rules and standards.

Acknowledgements We would like to thank Stijn Claessens, Sebastian Doerr and Jon Frost for comments and suggestions. We thank Giulio Cornelli for excellent research assistance. The views in this paper are those of the authors only and do not necessarily reflect those of the Bank for International Settlements.

References

Aaronson, S A. 2019, October. What are we talking about when we talk about digital protectionism? *World Trade Review* 18 (4).

¹⁷ For example, competition policy in the United States falls under the remit of both the Department of Justice and the Federal Trade Commission.

¹⁸ See <https://datagovhub.org/>.

- Acquisti, A, C. 2016. Taylor and L Wagman. The economics of privacy. *Journal of Economic Literature* 54 (2).
- Apple. 2016, February. A message to our customers. February 2016.
- Bar-Gill, O. 2019. Algorithmic price discrimination: When demand is a function of both preferences and (mis)perceptions. *University of Chicago Law Review* 86.
- Bureau Européen des Unions de Consommateurs (BEUC). 2019, November. Access to consumers' data in the digital economy. Position paper.
- BIS. 2019, June. Big tech in finance: Opportunities and risks. *Annual Economic Report 2019*, Chapter III.
- Carrière-Swallow, Y and V Haksar. 2019. The economics and implications of data: An integrated perspective. *IMF Departmental Papers* 19 (16).
- Coase, R. 1960. The problem of social cost. *Journal of Law and Economics* 3.
- Cyberspace Administration of China. 2019, June. Security assessment of cross-border transfer of personal information.
- Chen S, S Doerr, J Frost, L Gambacorta and H S Shin. 2020. The fintech gender gap, mimeo.
- EY. 2019. *Global FinTech Adoption Index*.
- Farboodi, M, R Mihet, T Philippon and L Veldkamp. 2019, January. Big data and firm dynamics. *NBER Working Papers*, no 25515.
- Financial Times. 2019a. How top health websites are sharing sensitive data with advertisers, 13 November, www.ft.com/content/0fbf4d8e-022b-11ea-be59-e49b2a136b8d.
- Financial Times. 2019b. The Indian payments conundrum, 18 November.
- Frost, J, L Gambacorta, Y Huang, H S Shin and P Zbinden. 2019, April. BigTech and the changing structure of financial intermediation. *BIS Working Papers*, no 779.
- Fuster, A, P Goldsmith-Pinkham, T Ramadorai and A Walther. 2019. The effect of machine learning on credit markets. *VoxEU*, 11 January.
- Galbraith, J. 1958. *The affluent society*.
- Gambacorta, L, Y Huang, H Qiu and J Wang. 2019. How do machine learning and non-traditional data affect credit scoring? Evidence from a Chinese fintech firm. *BIS Working Papers*, no 834.
- Hau, H, Y Huang, H Shan and Z Sheng. 2018. Fintech credit, financial inclusion and entrepreneurial growth, mimeo.
- Huang, Y, C Lin, Z Sheng and L Wei. 2018. Fintech credit and service quality, mimeo.
- Jullien, B. 2004, November. Two-sided markets and electronic intermediaries. *CESifo Working Papers*, no 1345.
- Kamer, A, J Guillory and J Hancock. 2014, March. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences of the United States of America* 111 (24).
- Mikians, J, L Gyarmati, V Erramilli and N Laoutaris. 2013. Crowd-assisted search for price discrimination in e-commerce: First results. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, 1–6.

- Miller, D, and P Friesen. 1984. A longitudinal study of corporate life cycle. *Management Science* 30 (10).
- Mitchell D and N Mishra. 2019, September. Regulating cross-border data flows in a data-driven world: how WTO law can contribute”, *Journal of International Economic Law* 22 (3).
- Nguyen Trieu, H. 2017. Why finance is becoming hyperscalable. *Disruptive Finance*, 24 July.
- Nilekani, N. 2018, September/October. Data to the people: India’s inclusive internet. *Foreign Affairs*.
- O’Neil, C. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*, Broadway Books.
- Petralia, K, T Philippon, T Rice and N Veron. 2019. Banking disrupted? Financial intermediation in an era of transformational technology. *Geneva Reports*, 22.

Index

A

- Aadhaar 838, 839, 841–849, 872
account-based 223, 224
account-based payments 139
account receivables finance 710
accredited investor 455, 491, 501
Acorn 27
Activation function 601, 627
Advertising 784, 786, 791, 795,
798–801
advisory role 464
AI R&D 632
AI research 595, 620, 621
Alcentra 719
Algebraic topology 631
Algorithmic aversion 742
algorithmic fairness 832
Algorithms/algorithmic decision-
making 784, 791, 794,
803
Alibaba 658, 659, 661, 666, 713
AliPay 13, 14
alternative currencies 95, 96, 100,
101, 103, 105, 121, 126,
130–132
Alternative finance 511, 512, 515,
517, 520–523, 526, 527,
532, 537–539, 541, 545–548,
550–552, 554–559, 567, 568,
590
alternative finance 191, 195, 196,
209, 211–213
Alternative Investment Fund (AIF)
485
Alternative Trading System 484
Amazon Lending 717, 722
Ambient space 599, 607, 631
American Bar Association (ABA) 329
analytics 671, 672, 675, 676,
678–680, 688, 694
AngelList 15, 16
Annuities 29
Ant Financial Services 712, 713
anti-money laundering (AML) 162,
167, 169, 227, 228, 307, 344,
459, 490, 492
App-based fintech lenders 210, 211
Application Programming Interfaces
(APIs) 752

area under the receiver operating characteristics. *See* AUROC

Article 148 463

Article 150 463

Artificial Intelligence (AI) 541, 546, 547, 550, 556, 566, 568, 590, 621, 671, 674–676, 681, 689, 693, 857, 873

Artificial neuron 600

A-share 480

asset management 733, 735, 737

asymmetric information 407, 412, 420, 856

Attention mechanisms 594, 637

AUROC 863, 864

Australia 465, 467, 468

Australian Competition and Consumer Commission (ACCC) 767, 768

Australian Securities and Investments Commission (ASIC) 465

Austria 465

automated loan 659, 666

B

backers 397–402, 404

Baihang Credit 816, 827, 833

bank accounts 840, 841, 846, 847

bank deposit insurance 447

benefits 841–843, 848, 849

bespoke hosted model (regulation of equity crowdfunding) 492–495, 498–500

Betterment 25

Bias term 598, 603

Bias-Variance Trade-off 598

big data 5–8, 11, 655, 656, 658, 659, 661, 662, 666, 671–673, 675, 681, 686, 690, 693, 695, 696, 856, 860, 861, 863, 864, 866

big data analytics 819, 822, 824, 826, 828, 832–834

Bitcoin 23, 24, 61, 62, 64, 69–77, 145, 243, 246–250, 256, 258–260, 262

Bitcoin Standard 156

Biz2Credit 715, 716, 722

Blending data sources 631, 632

blockchain 4, 12, 20, 27, 28, 64, 67, 69–77, 142, 243–250, 256, 257, 260–262, 269–272, 274–281, 283–285, 337–343, 345–364, 369, 370, 374, 377, 380, 382, 384, 385, 387, 388, 390, 671, 675, 692

Bloom 25

Blueprint Income 29

borrower protection 445, 446, 451, 453

Branch 197, 209–211

budgeting problem 738

business angel 477–480, 484, 492

C

Canada 465

capital flight 151

captive ecosystem 866, 868

cash 217–219, 223, 224, 226, 228, 237

Cashless Payment Systems 14

Cell state 613, 642

central bank digital currency (CBDC) 144, 162, 163, 169–174, 222

Central Bank of Brazil (BCB) 772

central bank public goods 224

central banks 51, 99, 103, 124, 125, 143, 144, 150–155, 157, 158, 162, 165, 166, 169–171, 173, 222, 224, 228, 236, 250, 416, 813, 873

Centralized digital currency 145

centralized identity 844, 845, 849

China 441–443, 450, 456–464, 468, 469

- China Banking and Insurance Regulatory Commission (CBIRC) 457, 459
- China Securities Regulatory Commission (CSRC) 459, 500, 501
- Chowdhry, Bhagwan 838, 844
- cloud computing 671, 692
- CMA9 757
- collateral 856, 860, 861, 865
- Commercial Bank of Africa (CBA) 194, 200, 201, 203
- company scrip 154
- competition 218, 235
- Competition and Markets Authority, UK 757
- compliance 689, 694, 695
- Computer vision 593–595, 616, 628
- Congress 451
- Connection weights 595
- consensus mechanisms 271, 281, 285
- consensus protocol 64, 74–76
- Consumer Data Right (Australian legislation) 767
- Consumer Financial Protection Bureau (CFPB) 452, 453, 796, 802
- consumer lending 34, 35, 54
- consumer protection 222, 232, 233, 451–454
- consumer welfare 869
- Convolution 607, 615, 616, 626
- Convolutional neural network 593, 595, 616, 617, 645
- correspondent banking 219–221, 236
- costs 841, 842, 844
- countering the financing of terrorism (CFT) 226, 228, 230, 233
- credit bureau 862, 863
- credit constraints 409, 420
- Credit Monitoring 12
- credit provision 658
- credit reporting 809–814, 816–819, 821, 823, 826, 827, 830–834
- credit scoring 810–814, 818–821, 823, 824, 826, 827, 859, 862–865, 867
- cross-border payments 219–221, 223, 228, 236, 237
- crowdfunding 4, 15, 397–404, 425–437, 444–448, 450, 464, 466, 467, 470
- Crowdfunding, Debt-Based 4
- Crowdfunding, Donation-Based 4
- Crowdfunding, Equity-Based 4, 17
- Crowdfunding Insurance 29
- Crowdfunding, Rewards-Based 4, 16
- Crowdfunding, Securities-Based 17
- crypto-assets 221–225, 245, 246, 249, 250, 256, 257, 260, 262, 263
- cryptocurrency(ies) 5, 20, 21, 23, 95–98, 101, 103–105, 113, 114, 116, 117, 120, 125–128, 130–133, 162–168, 170, 171, 175, 177–181, 338, 339, 341–343, 345–347, 351, 357, 364, 370, 371, 373, 376, 377, 388
- currency 221, 222, 226, 229, 231, 233
- currency digitization 142
- Curse of dimensionality 599
- custodian 462
- customer due diligence (CDD) 218
- customers 751, 753, 754, 756, 758, 760, 761, 763, 764, 767–770, 772
- Customized Portfolios 26
- D**
- Dalberg 846
- Data 844
- data 839, 843–845, 848, 849
- data leakages 746

data-network-activities. *See* DNA
 data protection
 authorities 873
 laws 871
 Data science 607
 data services 686, 687
 debt currency 99, 100, 115, 118,
 119, 127–129
 Decentralized Autonomous Organiza-
 tion (DAO) 70, 114, 177,
 302, 311, 312, 321, 325, 354
 decentralized digital currency 145
 Deep learning 589–596, 598–601,
 604, 605, 607, 609, 611, 620,
 622, 627, 629–632, 634–637,
 640
de minimis exception 327
 digital
 monopoly 866–868
 platforms 861, 868
 digital architecture 849
 digital cash 224
 digital currencies 142, 162–168,
 170–173, 176, 178, 180–182
 Digital Equivalent 146
 digital financial services 190, 192,
 194, 196, 212
 digital footprints 820
 digital identity 235, 237, 838,
 840–842, 848, 849
 digital money 221, 223, 225, 226
 Digital Parallel 146
 digital payments 218, 231, 234, 235,
 237
 digital revolution 190, 195
 digital signature 68, 71
 digital solutions 198
 Digital Wallets 22
 Dimensionality reduction 605, 607,
 637
 directed acyclic graph 76
 disclosure 479, 481, 490, 492, 497,
 503
 Discrimination 8, 10

Discriminator 628
 distributed ledger 64, 69, 71, 76, 77
 distributed ledger technologies 142,
 143, 162, 243, 256, 260, 262
 Distributed representations 590, 612
 diversification 479, 482, 502
 DNA 855, 856, 858, 859
 Dodd-Frank act 452
 Drew Hinkes 314
 Drèze, Jean 843
 due diligence 479, 481, 482, 484,
 486, 490, 492, 496, 498, 499,
 502–504
 Dukaton 143
 Dynamic memory 594

E
 early-stage firms 425, 428, 433, 434,
 436, 437
 early stage investor 479
 e-commerce 857, 859, 863, 865,
 868
 e-commerce platform 662, 664
 E-gCDC 144
 E-gDDC 144
 eGold 145
 1864 National Bank Act 451
 electronic-ID system 214
 electronic money (e-money) 218,
 221, 223, 224, 231, 232, 234,
 236, 237
 electronic payment networks 139
 Ellevest 26
 Elman networks 612
 Embedding 590, 605–607, 609,
 610, 612, 613, 617, 620,
 623–626, 631, 643
 emerging market and developing
 economies (EMDEs) 217–219,
 221–225, 227, 228, 230–234,
 237
 e-money 191

- Encoder-decoder model 591, 626, 637, 639
- end user 225
- entrepreneurial finance 426, 431
- entrepreneurship 391
- equity crowdfunding 477, 478, 480–483, 485–491, 496, 497, 500–502, 504
- equity finance 426, 434
- ETFs 22, 25, 26
- ethics 674, 689–693
- European Bank for Reconstruction and Development (EBRD) 490
- European Securities Market Authority (ESMA) 466, 484
- exchange 338, 340–343, 346–348, 350, 355–359, 363, 364
- exogenous references 153
- F**
- Facebook 795, 800, 801
- fast payments 235
- Federal Deposit Insurance Corporation (FDIC) 451–453
- Federal law 452
- Federal Trade Commission (FTC) 452
- Fiduciary Duties 463
- Filters 616
- Financial Access at Birth 840
- Financial Conduct Authority (FCA) 34, 417, 444–448, 450, 469, 478, 485, 497, 694, 757, 758, 763, 802
- financial exclusion 189, 196, 840, 843
- financial history 36, 41
- financial inclusion 193–196, 212, 213, 223, 226, 228, 237, 840, 844, 848, 856, 860
- financial integrity 222, 230, 232
- financial lending 409
- financial literacy 726, 728, 732, 737, 744
- Financial Planning 27
- financial regulation 451, 452
- financial regulators 873
- Financial Services Compensation Scheme 447, 449
- financial solutions 196, 199
- financial stability 869
- financial technology 471, 703, 721
- finTech 34–36, 41, 46, 55, 162, 166, 175, 176, 180, 197, 209, 210, 234–236, 238, 426, 483, 489, 511–513, 517, 518, 522, 523, 526, 528, 530–534, 536, 537, 540, 541, 546, 547, 555, 556, 559, 564–567, 655–664, 666, 704, 713, 715, 739, 744, 753, 760, 761, 763, 768, 785, 798
- France 464, 468
- fraud 9, 17, 20, 457, 463, 675–678, 676, 695
- Funding Circle 712, 719–721
- G**
- gatekeeper model (regulation of equity crowdfunding) 491, 498–501
- gCDC 144
- gDDC 144
- General Data Protection Law, Brazil 771
- General Data Protection Regulation (GDPR) 323, 767, 794, 800, 803, 873
- Generative adversarial networks (GAN) 627, 628
- Generator 628
- Germany 465, 466
- Glenshaw Glass Co.* 328
- Global Alternative Finance Regulation Survey 478
- GloVe 606

governance 288, 295–300, 303–308, 446
 government 838–841, 844, 845, 849
 Gradient descent 592–594, 597, 601, 604, 606
 GRU 595, 613, 615

H

hard forks 328
 hard information 659, 666
 Hebbian Theory 594
 herding 481, 482
 Hidden layer 593, 594, 603, 605, 606, 609, 616
 Hidden state vectors 624
 High dimensional data 610
 Hong Kong Monetary Authority (HKMA) 768–770
 hosted pitch model (equity crowdfunding) 481, 483, 484
 Howey Test 313
 Hyper-parameters 604, 635

I

iCDC 144
 iDDC 144
 identity 838–840, 842–845, 848–850
 Image captioning 619, 640
 incentives 840, 841, 843, 849
 inclusive growth 664, 666
 income aggregators 738
 India 467, 838–846, 849
 India Stack 873
 Indiegogo 16
 informality 218
 informational cascades 402–404
 information asymmetry 480
 Information Technology 838
 initial coin offerings (ICOs) 20, 21, 35, 42, 80, 85, 86, 122, 162, 163, 175–181, 288, 314,

369–376, 378, 381–391, 398, 402–404, 488, 489
 initial public offerings 42
 Innovative Finance Individual Savings Accounts (IFISAs) 444
 Input layer 606, 607, 616
 institutional investors 9, 455, 456, 469
 Instruction following 637, 640, 641
 insurance 28, 669–681, 683–687, 689–691, 694–697, 856, 859, 860, 866, 870
 InsurTech 28, 672, 675, 681, 683, 686, 687, 696
 Interim Measures 458–463
 Internal Revenue Service (IRS) 327
 internet 843, 844
 Internet of Things (IoT) 671, 673, 681, 682, 684, 686
 Investment Advisor Act of 1940 456
 investment managers 448
 Investment Platform Fragmentation 21
 investor protection 450, 454, 458, 463, 465, 467–470
 invoice finance 702, 710, 719

K

Kabbage 7, 11, 715, 721
 Kenya 141
 Kickstarter 16
 Kiva 4, 8, 9
 know your customer (KYC) 28, 226, 228, 235, 708

L

Language grounding 639
 legal 443, 451, 453, 454, 459, 462, 466
 Lemonade 30
 lending 856, 859–861, 864
 LendingClub 6, 7, 452–455, 469

lending limits 462
 lending platforms 441–448, 450,
 451, 453–459, 462, 464–471
Lex Cryptographia 326
 Libra 156
 life cycle 856–858
 Life Insurance 29, 30
 Linear structure 594, 616
 loan-based crowdfunding 445, 446,
 448, 450, 464, 466, 470, 703,
 705
 loans 859, 861, 863–865
 local and complementary currencies
 95, 105, 111, 126
 Local Currency 24
 Local Investors 18
 LocalStake 18
 Local Ventures 18
 Loss function 594, 601, 603, 604,
 644
 LSTM 590, 595, 612, 613, 624,
 627, 630, 639, 643

M

Machine learning 591, 592, 595,
 597, 598, 600, 609, 630, 635,
 639
 machine learning 857, 861–865, 867
 Madden v Midland Funding LLC
 451, 454
 Manifold 589, 596, 599, 607, 631
 Manifold theory 627
 margins 858
 market frictions 656
 MarketInvoice 718, 719
 marketplace 442, 443, 450, 452,
 453, 455, 464, 469
 marketplace lending 442, 443, 453,
 469
 market power 856, 866, 867
 microfinance 34, 36, 51–55
 Microfinance Institutions (MFIs) 8,

9

Microlending 8
 Mint 27
 mobile banking 237
 mobile money 222, 228, 231, 234,
 235
 mobile network operators (MNOs)
 194, 197, 213
 mobile payment 660, 661, 663, 664
 mobile phones 846, 848
 mobile phone transactions platform
 189
 mobile virtual network operations
 (MVNOs) 194, 197, 211, 213
 Modi, Narendra 841
 Modular Open Source Identity
 Platform (MOSIP) 849
 MoKash 194, 200–203, 205, 206
 MoMoKash 200, 203–206, 209
 Monetary Authority of Singapore
 (MAS) 770, 771
 monetary policy 151
 money 95–97, 99–104, 107–110,
 112, 115–122, 124–133,
 217–219, 222, 224, 225, 228,
 234
 money management 856, 859
 money transfer operators (MTOs)
 219
 Monitoring 7–10, 12, 14, 16, 28, 29
 monitoring 408, 418–420, 861, 865
 M-Pawa 194, 200, 201, 205, 206,
 209
 M-Pesa 141, 189–194, 197, 200,
 205, 210, 211
 M-Shwari 192–194, 197, 200,
 204–206, 208, 211–213
 MSP 857, 858
 Multilateral Trading Facility 484
 multi-sided platform. *See* MSP
 multi-stakeholder collaboration 308
 MYbank 713, 714
 MyInfo 872

N

National Banking and Securities
 Commission, Mexico 774

Natural language understanding 637

Netherlands 464

network
 analysis 861, 862
 effects 84, 86, 855, 857, 858, 865
 externalities 855–858, 866

Neural activity 631

Neural architecture search 629, 630

Neural machine translation 617

Neural networks 591–594, 596, 599,
 601, 604, 611–613, 616, 617,
 628, 631, 632, 635, 636

Neuroscience 592–596, 599, 617,
 627, 633

New Venture Financing 14

New Zealand 464, 467

NextSeed 18

N-gram 598

Nilekani, Nandan 838, 842

1933 ‘Truth in Securities Act’ 452

NLP 590, 594, 595, 597, 600, 605,
 607, 617, 620, 622, 623, 626,
 634, 636

nonce 71, 72

Non-Mainstream Pooled Investment
 485, 491

non-performing loans (NPLs) 203,
 204, 207

non-recourse notes 454

O

Office of Comptroller of the
 Currency (OCC) 452, 453

Office of Fair Trading (OFT) 445

Office of Fair Trading to the
 Financial Services Authority
 444

Office of the Australian Information
 Commissioner (OAIC) 767

Ondeck 714, 715, 721

online credit markets 409–411, 414,
 416, 419, 420

Open Banking 235, 237, 722, 872

Open Banking Implementation
 Entity (OBIE) 757, 758, 761,
 763, 764

Open Banking Working Group (UK)
 756, 775

Open Finance 759, 774–777

Output layer 602, 609, 644

Overfitting 598, 600

ownership of data 869

P

P2P 442, 444–446, 449, 450,
 456–463, 467–469

P2P platforms 457, 459, 460, 463

Parameters 590–592, 595, 598, 599,
 611, 616, 644

payments 79–81, 83, 84, 88, 89, 91,
 856, 857, 859, 860, 865, 872
 payment services 859, 860

Payment Systems 13

Paypal 718

PBOC 457, 461

peer-to-peer 470

Peer-to-Peer Insurance 29

Peer-to-Peer Lending 9

peer-to-peer market 411, 416, 417

peer-to-peer platforms 411, 416,
 417, 789

People’s Bank Of China. *See* PBOC

Perceptron 593

permissionless 273, 275, 276, 278,
 280, 282, 283, 285

permissionless innovation 288, 294,
 299

personal data 856, 857, 865,
 869–873
 personal transaction data 861

Personal Finance 21

personalised services 675, 679, 681

Petro 156

- P-gCDC 144
 - P-gDDC 144
 - photo identity 838, 845
 - platforms 86, 441–451, 454–471
 - Poland 465
 - political 841, 842
 - Pooling 616
 - portable identity 839
 - poverty 840, 841, 843, 848
 - pre-emption 452, 453
 - Preprocessing 605, 609
 - price discrimination 7, 29, 856, 866–868
 - privacy 5, 13, 29, 813, 816, 827, 830–832, 838, 844, 845, 849, 856, 869–871, 873
 - Private Benefits 19
 - private currency 46–50, 55
 - private key 68, 69, 71
 - programmable money 152
 - Project Jasper 143
 - Project Khokha 143
 - Project Stella 143
 - Project Ubin 143
 - proof of stake 75, 271, 281, 283, 285
 - proof of work 74, 271, 273–277, 280, 281, 283, 285
 - prospectus 497, 499
 - Prosper 6, 7, 10, 450, 453–455, 469
 - Protected Information 8
 - Proxy variables 791, 797
 - Psychometrics 532–534, 540, 541, 543, 544, 546, 547, 550, 556, 558, 564, 566, 568, 596
 - Public accommodation laws 785, 795, 803
 - public goods 224
 - public infrastructure 844
 - public key 68, 69
 - public offering model (regulation of equity crowdfunding) 490, 493–495, 497, 500–502
 - public policy 856, 869
 - purchase order finance 702, 710, 719
 - Python 633, 636, 644
 - PyTorch 633, 635–637, 644
- Q
- QR code 663
 - QuantmRE 26
 - Qudian 6
- R
- Ratesetter 444
 - ration card 843
 - Ravencoin (RVN) project 330
 - Real Estate 26
 - Real Estate Model (Equity Crowdfunding) 483
 - real time gross settlement (RTGS) 235
 - Reconstruction 609
 - Record Keeping 12
 - Recurrent cell 590, 610, 632
 - Recurrent neural network 611, 612, 619, 635, 641
 - redistributive effects 664
 - RegTech 471, 694, 695
 - Regularization 594, 616, 642
 - regulation 7, 10, 17, 20, 34, 35, 37, 39–41, 43–46, 49, 51, 53–55, 131, 152, 158, 162–164, 166, 167, 175–178, 181, 207, 224, 233, 245, 251, 252, 254, 258, 259, 289–291, 293–295, 300–306, 353, 360, 361, 436, 441–446, 448, 450, 451, 453, 454, 458, 463, 465–470, 484, 488–490, 502, 691, 692, 694, 695, 761–765, 768, 770, 772, 774, 777, 785, 786, 793, 803, 864, 865, 873
 - regulatory policy 864
 - Regulation A 315

Regulation CF 316
 Regulation D 315, 455, 485
 regulatory 443–446, 450, 451, 453, 457–460, 464–471
 regulatory frameworks 232, 233
 Reinforcement learning 590, 594, 597, 600, 628, 629
 remittances 219, 221, 226–229, 234, 237
 reregulation 469, 470, 722
 retail electronic payments platform 189, 191, 193, 212, 213
 retail investor protection 467, 470
 Revenue-Share Agreements 19
 reward-based 397, 399–402, 404
 risk 670, 672, 674, 675, 678, 682–685, 688–690, 694, 696, 697
 RNN 595, 611–613, 630, 639, 643
 robo-advising 725–730, 732–737, 739–746
 Robo-Advisors 5, 24
 Robotics 639

S

Safaricom 194, 200
 Sardex 24
 screening 409, 412, 413, 861
 search engines 857, 859, 868
 Second Circuit 451, 454
 Second Payment Services Directive (PSD2) 755–758, 764, 765
 Securities Act of 1933 312
 Securities and Exchange Commission (SEC) 450, 454, 455
 securities-based 402, 403
 SeedInvest 15, 16, 18
 Sesame Credit 821, 826, 828, 829, 832
 SHA-256 66–68, 71
 shadow banking 864
 Silk Road 324
 Similarity 605, 606, 610

Simple Agreement for Future Equity (SAFE) 19
 Simulation 600, 604, 628
 Singh, Manmohan 841
 Slovakia 465
 Slovenia 465
 small and medium size enterprises (SMEs) 701, 703–707, 709, 712, 713, 715, 717, 719, 721
 smart card 843
 smart contract 318
 Social Credit System 149
 social networks 6, 857, 869
 social relation 100, 101, 127–129, 132
 Social Security Number 838
 SoFi Wealth 25
 soft information 4, 11, 862, 863
 Sophisticated Investors 15
 Spain 464
 stablecoins 152, 221–234, 237
 StartEngine 18
 Statistical discrimination 784, 787, 788, 791–794, 796, 798, 803
 Stock market prediction 591
 subsidies 843
 Sweden 140
 Syndicated/Hybrid Crowd model (Equity crowdfunding) 483, 484, 492
 Synthetic data 600
 Szabo, Nick 318

T

Tala 197, 209–211
 tamper-resistance 146
 Target Date Funds 728
 Taste-based discrimination 787, 788, 803
 tax-loss harvesting 736
 technology 842, 844, 849
 Telco-facilitated bank virtual savings and credit products 209

- Temporal dependency 610, 613
 Tensor computation frameworks 633
 TensorFlow 630, 633, 635, 636
 Tether 152
 Third Party Provider (TPP) 753
 Time series 590, 605, 609–611
 Timicoin 28
 token 369–378, 381–391
 token-based 223, 224
 token economy 246, 260, 261
 tokenization 341–343, 345, 346
 trade finance 702, 711, 712, 720, 721
 Training data 596, 599, 601, 628, 631, 639, 644
 Transfer function 603
 true lender 453
 trust building 810, 811, 817, 818, 821, 825, 828–830, 833, 834
 Truth in Securities Act 452, 455
 2006 Chinese Company Law 463
 2006 Company Law 463
 2015 Guiding Opinions 459
 2017 Guideline on Information Disclosure 458, 462, 463
 TZero 27
- U**
 UK 703, 705–709, 712, 717–719, 721, 722
 unbanked 141, 171, 194, 209, 212, 231, 234, 856, 859, 860, 862
 Underfitting 598
 Unified Payment Interface. *See* UPI
 Unique Identification Authority of India (UIDAI) 839, 841, 844, 845
 Unregulated Collective Investment Schemes (UCIS) 485
 unsecured credit 710
 Unstable gradient problem 613
 UPI 872
- US 441–443, 450, 451, 453–456, 466, 468, 469, 705–708, 712, 714–717, 719, 721
 usury laws 453
 utopia 96, 100, 102–105, 107, 108, 112, 113, 120, 126, 127, 130–133
- V**
 Value-at-Risk 736
 value(s) 96, 101, 108–117, 119–122, 124–126, 129, 131, 132
 Vapnik–Chervonenkis Theory 597
 Vector spaces 607
 Venezuela 156
 Venmo 13
 virtual credit supply platform 189, 194, 210, 212
 virtual platform 194, 195, 212
 virtual savings and credit products 195, 197, 200, 204, 205, 208–210
 virtual savings platform 192
- W**
 Wealthsimple 25
 WeCash 6
 WeChat 13, 14
 Wefunder 18
 welfare 841, 846–848
 WhopperCoin 154
 wisdom of crowds 398, 399, 403, 404
 Word2vec 606, 607, 624, 626
 working capital finance 701–704, 706–710, 712–714, 721, 722
 World Bank 477, 478
- Y**
 YieldStreet 21

Z

Zelle [13](#)

Zero-Lower-Bound [151](#)

Zopa [443](#)