The New World of Blockchain Economics: Consensus Mechanism as a Core Element

Jona Stinner and Marcel Tyrell

1 Introduction

Roughly a month following Lehmann Brothers' bankruptcy at the peak of the financial crisis in autumn 2008, a brief article, titled "Bitcoin: A Peer-to-Peer Electronic Cash System", appeared in a cryptographic mailing list. Its pseudonymous creator, Satoshi Nakamoto, envisioned an electronic cash network as an alternative to the traditional financial system, which at that time faced existential distress (Nakamoto, 2008). Inconspicuous at first glance, the whitepaper paved the road to the first cryptocurrency "Bitcoin", presently valued with ≈$40,000/BTC (January 2021) and a market capitalization of nearly $1 trillion.[1]

Bitcoin is an explicitly decentralized, digital peer-to-peer network, which allows participants to exchange financial value in the form of virtual currency tokens (Nakamoto, 2008). Such tokens combine the benefits of a digital settlement infrastructure with some qualities commonly associated with cash, including third-party anonymity, quick payment finality, small transaction costs and mutual exchange between parties, however, and that is important, without the (direct) involvement of any intermediary as trusted third party.[2]

Transferring these properties to a purely digital currency has historically been challenging, as illustrated by the juxtaposition with cash: First, cash transactions

[1] Data from https://cryptoslate.com/coins/, accessed in January 2021.

[2] See Narayanan et al. (2016) for one of the earliest comprehensive descriptions of PoW-consensus.

J. Stinner (✉) · M. Tyrell
Department of Philosophy, Politics, and Economics, ISIC—The PPE Institute for Social and Institutional Change, Witten/Herdecke University, Alfred-Herrhausen-Str. 50, 58448 Witten, Germany
e-mail: jona.stinner@uni-wh.de

M. Tyrell
e-mail: marcel.tyrell@uni-wh.de

T. Herberger and J. Dötsch (eds.), *Digitalization, Digital Transformation and Sustainability in the Global Economy*, Springer Proceedings in Business and Economics, https://doi.org/10.1007/978-3-030-77340-3_2

can be easily processed and verified between any physically present individuals. Instead, digital accounts are at constant risk of fraudulent access (i.e., cyber-attacks), requiring enhanced validation processes. Moreover, digital processing needs a mediating provider, who adjusts the accounts according to the initialized transfer and manages the underlying infrastructure. Second, digital settlement often induces a delay between goods provision and payment finality.[3] This delay exposes the seller to risks if payment reversibility is a feasible option, whereas cash transactions guarantee conclusive compensation. Third, electronic files can be duplicated effortlessly and at negligible costs. If unrestricted, this allows individuals to inflate the currency at will and spend amounts multiple times (referred to as "double-spending-problem") (Berentsen & Schär, 2018). While central banks have developed anti-counterfeit features to prove the authenticity of cash, guaranteeing this quality for electronic money is a non-trivial exercise (Deutsche Bundesbank, 2017).

The advent of modern asymmetric cryptography provided solutions to validate the legitimacy of digital communication (e.g., Hellman, 2002; Naor & Yung, 1989). Providers of virtual payment systems employed those innovations to ensure reliable account access and validate transfers, thus overcoming the first described obstacle of a trusted validation. But to counteract the double-spending problem, earlier implementations of electronic cash[4] and conventional monetary and payment systems had to rely on a central intermediary as trusted third party. However, as became apparent in the cause of the global financial crisis (i.e., the collapse of Lehman Brothers 2008), centralized systemic relevance of a single or profoundly connected set of entities can cause undesirable externalities and substantial problems.

In contrast, Bitcoin bypasses the requirement for a trusted third party by (i) removing the autarchy of a central record keeper in favor of a publicly shared ledger; and (ii) shifting system stewardship to a set of anonymous and widely distributed nodes. Record-keeping and transaction settlement in the Bitcoin network is built on a novel digital infrastructure, termed "blockchain". A blockchain is a cryptographically secured, distributed ledger that is copied a thousand-fold across any network member willing to download and operate it. Each copy of the blockchain stores a complete history of transactions in an interlocked chain of blocks, according to the rules of an underling protocol (Nakamoto, 2008).

Some nodes, called "miners", perform the process of altering the record when new transactions are initialized in the network. Miners verify and bundle pending transactions into block candidates in a cryptographically robust way and broadcast them to the network. Other nodes express their acceptance of a valid candidate, by appending it to their version of the ledger and continuing on the expanded sequence (Böhme et al., 2015). This process allows the network to exchange transactions in a continuously enlarging ledger. However, to provide a solution for the double-spending problem (and obtain a genuinely decentralized ecosystem) the network must consent upon a unique state of ownership among the arbitrarily distributed participants and numerous copies of the ledger. Nakamoto's core innovation is the

[3] In particular, if multiple entities are involved such as in an international money transfer.

[4] E.g., DigiCash or e-gold.

incentive mechanism under which this decentralized process operates (more or less) frictionless.

Bitcoin's consensus protocol replaces trust in a central entity's honesty with cryptographic principles and economic incentives, designed to align participants' selfish economic interests with collectively intended behavior. Such a protocol is fundamental to the architecture of each permissionless cryptocurrency, but some advancements have been made regarding its specific technical and economical design. In fact, the invention of Bitcoin initiated the formation of a universe of currently $\approx$2500 cryptocurrencies with significant proliferation and substantial heterogeneity in the approaches to obtain consensus (Irresberger et al., 2020).

This article aims to illustrate the fundamental structures of two of the prevailing consensus protocols, namely proof-of-work (PoW) and proof-of-stake (PoS), and to highlight some important distinctions. In particular, it describes how the rules stated in each type of protocol enable the members of a self-sufficient cryptocurrency to coordinate to and agree on a unique state of ownership. Section 2 describes proof-of-work, which gained considerable attention as the first functional consensus protocol in a decentralized ledger technology. Since their evolution is closely related, Sect. 2 includes Bitcoin as an example of an applied PoW mechanism. Section 3 continues with a short description of one widely discussed alternative, the so-called proof-of-stake mechanism. In Sect. 4 we point to some nonstandard consensus protocols, particularly the group of committee-based-consensus/Byzantine Fault Tolerance mechanism and provide concluding remarks.

2 Proof-of-Work Consensus

A consensus protocol consists of a set of rules that, when applied collectively, formalize the conditions under which new transactions are attached to the ledger of a decentralized payment system. By design, network participants are strongly encouraged to follow the specified rules, as conformity maximizes individual gains and is thus in their best interest (Huberman et al., 2017). To understand how the PoW mechanism of a blockchain creates such a precondition, we examine its constituents with regard to the most prominent example, the Bitcoin network.

The transfer of virtual currency in the Bitcoin network is organized as follows: A participant willing to conduct a transaction broadcasts to the network that another wallet address should be the new owner of a specified amount of Bitcoins. Public-key encryption guarantees the sender's legitimacy and integrity of the encrypted transfer information when the initialized order is distributed among the network members. Slightly simplified, the underlying principle is based on a cryptographic pair of keys. The public key identifies a specific Bitcoin address (i.e., comparable to an account number) and is apparent to all network members. The owner of the address keeps the private key secretly (i.e., analogous to an account password). When initializing a transfer, the owner uses the private key to assign a digital signature to the transaction message before spreading it across the network. Any network member is able to decipher the signature with the respective public key and inspect or alter the contained information at will. However, restoring the signature requires the private key, which is solely known by the owner, rendering the message invalid in case of (malicious) modifications. Hence, a valid signature proves both the sender's authenticity and the immutability of the contained information (Schär & Berentsen, 2020).

To append transactions to the blockchain, miners collect pending messages, verify the signature's validity, and include them into a block candidate. In order to become accepted by other peers (and subsequently be included in their version of the distributed ledger), block candidates must (i) contain only valid transactions and (ii) correspond to the target of a difficulty criterion, as explained shortly. The standardized integration of public-key authentication makes it straightforward for miners to meet the first condition. The second condition, however, intentionally provokes substantial effort in the form of repeatedly operated hash computations. The devoted work and associated costs are essential for maintaining the network's trustworthiness and security. However, before examining this implication in more detail, the hashing operation has to be described in more detail.

A cryptographic hash is the output of a one-way function that transforms an arbitrary length input into a fixed-length value.[5] Some properties of hash values make them particularly useful for the application in cryptocurrencies: First, a hash value uniquely identifies the input (somewhat analogous to a digital fingerprint). Even infinitesimal modifications of the input lead to a substantially and unpredictably different hash value. Second, hash values are drawn from an astronomically large distribution of possible manifestations. Hence, it is stochastically implausible to receive a colliding value for non-identical inputs. Third, hash functions are non-invertible, i.e., it is infeasible to reconstruct the input if only the hash is visible. Conversely, due to the quasi-random transformation of the function, it is impossible to selectively customize the input to obtain a desired output (Al-kuwari et al., 2011).

A critical application of hash functions in the Bitcoin ecosystem is to guarantee immutability of past transactions. Miners propose a block candidate by hashing its entailed transactions together with the hash value of the preceding block. Since each block is interlocked with its predecessor, this forms a tamper-evident chain of blocks, i.e., a blockchain. Altering any contained information implies changing the hash values across the entire chain and being notified immediately. However, multiple branches of the ledger may emerge, if, for instance, miners integrate different transactions from the pool of pending ones. That means, different block candidates will be formed (Biais et al., 2019). Therefore, the protocol restricts the generally permitted solutions by accepting block candidates only if the corresponding hash value starts with a certain number of zeros. As we will describe in the next paragraph, the number of starting zeros expresses the difficulty of generating the required hash

[5] Hash functions exist in a variety of specifications (e.g., SHA-1, mumur32 or MD5). Bitcoin is based on the hash function SHA256, which assigns a hash value of 64 digits for any practically length of input. For Instance, the SHA256-hash of the term "Bitcoin" results in "346f1bd1c9f195e21e81d662123c3c94c9df4fa728223459216f8e3a19862c2d", while the term "bitcoin" (notice the small deviation) gives "d67509e37480c61b46c6bfc26f66f6a1884614763f08db3ab45616c8c3c12983".

value. Given the quasi-random nature of hash functions and the current technology level, brute-force[6] is the exclusive strategy to find such a suitable hash value for appending a block to the Bitcoin blockchain. In practice, miners repeatedly adjust a random number within the block candidate, called "nonce", and compute the hash function until a value appears that fits the hurdle rate set by the blockchain algorithm. Of course, computational power (and a bit of luck) is the primary determinant of success (Schär & Berentsen, 2020).

Adding a "difficulty" to the puzzle (e.g., in the form of required zeros) ideally should lead to the result that a single miner (presumably the fastest) first broadcasts a valid candidate to the network. Nevertheless, multiple valid versions of the blockchain may occur, if (i) different miners simultaneously find a suitable hash value (which is statistically unlikely but happens occasionally) or some network members decide—for unspecified reasons—not to follow the first proposed solution. Preventing perpetual disagreement in this context requires a fundamental consensus rule: For the Bitcoin network, participants consider the most extended block sequence as valid. Forks (i.e., valid alternative branches of the ledger) may still exist temporarily if simultaneously proposed candidates result in multiple chains of equal length. However, since block creation follows a stochastic process, one branch necessarily exceeds the others sooner or later and is subsequently adopted (Halaburda et al., 2021).

Considering the longest chain as valid is clearly beneficial for each participant, who expects other network members to behave accordingly. Obviously, any potential benefits from working on an incompatible blockchain cannot be utilized and are therefore worthless, while its operation requires costly computational power. By proposing a candidate with a generally accepted hash, miners prove an externally verifiable commitment to the blockchain's current state. This fundamental principle of proof-of-work makes a fraud endeavor prohibitively expensive. For instance, without mining costs, miners may create undisclosed parallel chains in order to double-spend currency. If such a private chain is provided with valid hash values, it is adopted by the network as soon as it exceeds the publicly shared blockchain, following the longest chain rule (Savolainen & Ruiz-Ogarrio, 2020). Suppose an attacker with enough computational power to outpace the honest network uses her currency on the main chain to purchase goods or services. After receiving the benefit, she publishes a secretly mined private chain, which is accepted by the network (Eyal & Sirer, 2014). The attacker makes a profit since she never moved her funds on the private chain, which allows spending the same amount of currency again (Budish, 2018).

Such an attempt becomes exponentially difficult for any miner controlling less than 50% of the total computational resources. Nevertheless, when abstracting from mining costs, expected returns of a double-spending-attack remain marginally positive even for smaller stakes. Establishing mining costs in the form of proof-of-work,

[6] I.e., the enumeration of potential candidates until a solution is found.

however, dramatically changes this simple calculation to the attacker's disadvantage.[7] A fraudulent miner then has to weight the uncertain return from an attack against its costs and the profits from honest behavior. Moreover, accumulating the computing resources required to conduct the described attack with a sufficient probability creates itself a substantial and irreversible financial burden for any potential attacker. The essential economic principle to note here is that the blockchain's immutability increases with the expenditures incurred for mining (Budish, 2018; Ma et al., 2018).

To attract a sufficiently large number of contributors willing to incur those costs, the protocol allocates a reward of currently 6.25 BTC and any attached transactions fees to the miner, who first appends a generally accepted block to the ledger. Reward levels bisect approximately every four years (more specifically, every 210,000 blocks) from originally 50 BTC to ultimately zero with a limited total supply of 21 million BTC. Since mining is permissionless (i.e., anyone can deploy computing power to mine blocks without any special permission), its market structure is characterized by pure competition. It follows that the presence of rewards encourages market entry until expected returns equal the incurring cost. Since only the winning miner (pool) is rewarded, the mining process is akin to a repeatedly executed tournament with free entry.

The incentive compatibility of PoW inevitable demands such a computational contest. However, this consensus design contains some obvious drawbacks: Bitcoin exhibits low scalability, as analyzed thoroughly by John et al. (2020), and requires substantial amounts of resources (de Vries, 2018). For instance, Bitcoin's aggregated energy consumption for the year 2021 is estimated at around 111.3 TWh—exceeding countries such as the Netherlands or Austria (Rauchs et al., 2021). For this reason, several recent projects in computer science have emerged to invent other consensus mechanisms that are more cost-effective and scalable.

3 Proof-of-Stake Consensus

One of the early developed alternatives, and nowadays still the main contender of PoW is the Proof-of-Stake (PoS) consensus protocol, which for example is used by Cardano and Binance Coin (Irresberger et al., 2020).[8] The fundamental principle of the PoS protocol is the following: Instead of relying on competitive forces, as implemented in the PoW protocol, the PoS mechanism authorizes a randomly selected stakeholder, called "validator",[9] to update the ledger (Saleh, 2020). This means that

[7] In particular, fraudulent miners have to take into consideration that the attack probably has an adverse impact on the market value of the cryptocurrency.

[8] Irresberger et al. (2020) document a growth in number and relevance of the PoS blockchains, both in absolute numbers and also relative to PoW blockchains.

[9] Nodes, which are selected in the PoS system to update the ledger are also called "validators" since their task primarily consists of validating transactions, whereas PoW consensus requires nodes to repeatedly execute complex hash-computations.

no incentives exist for validators to engage in a computational contest, which—at least in principle—for the following reason could threaten or undermine the ability of the PoS design to generate consensus: Even though the PoS mechanism specifies that a validator is offered an explicit monetary reward,[10] no explicit monetary cost must be burdened to gain the authority to update the blockchain. Therefore, from a design perspective, the PoS consensus mechanism does not couple the benefit of earning the block reward with an explicit cost component—in contrast to the PoW mechanism. Critics argue that this creates incentives for the validator to always add a block to the chain, irrespective of whether the update induces perpetual disagreement among users. This is what these detractors call the Nothing-at-Stake problem. Of course, if true, that would destroy the PoS's long-run viability because in that instance an initial disagreement would persist indefinitely (Halaburda et al., 2021; Saleh, 2020).

This issue is theoretically analyzed by Saleh (2020), who, however, concludes that the Nothing-at-Stake problem is not valid in general. He makes the following reasoning: One key component of a PoS blockchain is that such a blockchain possesses a native coin (token) that must be used to facilitate exchange on that blockchain. Therefore, a stakeholder of a given blockchain refers to the native coin holdings of that individual. The PoS algorithm will consider nodes to append to the blockchain based on the number of tokens they put at stake. Thus, each node will be selected based on a probability distribution, which depends on the respective amount of tokens staked. The stakeholder who has most "on stake" therefore also has the highest probability of being selected to append a block candidate to the blockchain. The selected node adds the next block and receives the sum of rewards and transaction fees of the transactions comprising that block (Halaburda et al. 2021).

Saleh (2020) argues that proponents of the Nothing-at-Stake problem implicitly assume price-taking behavior by a validator, i.e., the validator does not take into consideration the effect her decisions will have on the native coin value of that blockchain. However, this seems not to be an appropriate behavioral assumption for a validator: If the validator chooses to update the blockchain in a way that perpetuates disagreement among nodes, she imposes a cost on all stakeholders. Such an action undermines users' ability to exchange the native coin since the coins' ultimate ownership history cannot be clearly determined. The validator creates a fork. As a result, the value of the native coin will decrease, given that its usage as an exchange medium is clearly hampered. Note that by taking such an action, a validator also imposes at cost on herself. Therefore, if she adds a block to the blockchain in a manner that leads to persistent disagreement, she herself has to incur a financial cost, which in term of magnitude depends on her stake in the blockchain. Thus, the stake might signal a validator's intent to behave honestly. This means that not only for the PoW but also for the PoS consensus mechanism, a financial cost is coupled with the benefit of grabbing the block reward and the transaction fees.[11]

[10] Similar in this respect to PoW consensus.

[11] For a similar argument in PoW consensus see Budish (2018).

A second important implication follows from this argument (Saleh, 2020): By restricting the access to the blockchain to sufficiently large stakeholders, an equilibrium can be induced that generates consensus (Fanti et al., 2019). As previously mentioned, the cost of updating the blockchain in a way that causes persistent disagreement increases with a validator's stake. A sufficiently high minimum stake requirement as a condition for gaining access to the validator function creates a situation in which the cost of persistent disagreement outweighs the benefits from adding "improper" blocks for all validators. The result is an equilibrium, in which all participants coordinate when updating the ledger to generate consensus (Saleh, 2020). Saleh concludes that using a PoS protocol for a blockchain not only avoids prohibitive energy consumption but also requires only a modest reward schedule, without endangering the viability of the consensus mechanism.

Of course, there is an ongoing discussion about whether the analysis of Saleh (2020) allows an appropriate implementation of PoS consensus. In the interest of brevity, we do not provide an overview here but refer in the footnote to the existing literature.[12] Instead, we want to point to one very interesting economic implication: Does the consensus protocol matters for scaling blockchains? John et al. (2020) investigate this question. On the basis of a model analysis, they indicate that it depends on the consensus protocol, i.e., either PoW or PoS, whether increasing the transaction rate has positive or negative effects on the security of the respective blockchain. In general, increasing the transaction rate (e.g., by extending the block size to process more transactions) mitigates congestion, resulting in decreasing transaction fees (Easley et al., 2019). The reason is that agents in equilibrium have to pay only proportionally smaller fees to obtain a comparable settlement speed when the processing capacity increases. In PoW blockchains, validators will consequently withdraw computational power from the network, as dictated by the incentive compatibility. In turn, a reduction of the computational power operated in the network lowers the cost of a successful attack (Budish, 2018) and therefore undermines the security of a PoW blockchain. It follows that increasing the scale of a PoW blockchain has a negative effect on network security.

In contrast, lower congestion leads to higher demand in a blockchain operated via a PoS protocol, which increases the market value of the native coin, i.e., the blockchain's cryptocurrency. This, in turn, increases the opportunity costs of an attack, given that the value of a validator's stake increases. Therefore, increasing the scale of a PoS blockchain improves its security (John et al., 2020). The results show that the design features of the PoS consensus mechanism are qualitatively very different from the PoW protocol. The computational power expended by validators of a PoS blockchain has no effect on its security. This feature makes the PoS consensus mechanism a particularly interesting alternative.

[12] E.g., Gans and Gandal (2019) and Halaburda et al. (2021). Savolainen and Ruiz-Ogarrio (2020) provide an interesting discussion of miners' incentives to attack and double spend in PoW blockchains, which is related to the concept of endogenous coin value in PoS consensus. By analyzing the occurrence of large mining pools, they argue that static costs (mining) are less important compared to dynamic incentives (revenue flows) to maintain honesty.

4 Conclusion

This paper gives an overview of the block building process and principles of different consensus protocols. We focused our descriptive analysis on the consensus protocol, which sets the rules for updating a given blockchain. Since consensus, i.e., the agreement on a unique ledger content, is key to the functionality of a decentralized record-keeping system, a more profound knowledge of these protocols provides a thorough understanding of the comparative advantages of different blockchain designs. Therefore, we analyzed the two most prominent economic protocols PoW and PoS, concerning the mechanisms under which the respective network achieves consensus. Of course, the public blockchain ecosystem consists also of many other alternatives, for instance, blockchains using hybrid PoW/PoS protocols and so-called Delegated PoS blockchains.[13] Since research on new consensus protocols is an active field, numerous other blockchain protocols were proposed.

A particularly interesting group are Byzantine consensus-based blockchains, such as Honeybadger, HotStuff or Tendermint.[14] Moreover, the Diem project proposed by Facebook, which claims "to build a trusted and innovative financial network that empowers people and businesses around the world", is based on a blockchain design using the Byzantine Fault Tolerant (BFT) consensus approach.[15] Amoussou-Guenou et al. (2020) offer a methodology to analyze Byzantine consensus-based blockchain protocols from a game-theoretic perspective. The protocols are set up as a committee coordination game between rational and Byzantine players. The authors specify conditions under which consensus properties are satisfied and show that coordination failures could be a big problem.

Trust in the record-keeper is the mechanism to establish a consensus in a traditional intermediary based centralized record-keeping system. As a fascinating decentralized alternative, the blockchain could revolutionize our traditional financial system. However, to live up to these expectations, further research on the viability of different consensus mechanisms is required. The results described in this paper prove that it is a worthwhile endeavor.

References

Al-kuwari, S., Davenport, J. H., & Bradford, R. J. (2011). Cryptographic hash functions: Recent design trends and security notions. *Eprint.Iacr.Org*, 1–36. Retrieved from http://eprint.iacr.org/2011/565.pdf

[13] Irresberger et al. (2020) provide an overview.

[14] A Byzantine fault tolerant system takes into account false messages provided by Byzantine players and allows for temporary disagreements among users (Halaburda et al., 2021). For more information on the practical implementation, see the whitepapers of Honeybadger, HotStuff and Tendermint.

[15] See Diem project homepage.

Amoussou-Guenou, Y., Biais, B., Potop-Butucaru, M., & Tucci-Piergiovanni, S. (2020). Rational behavior in committee-based blockchains. *IACR Cryptology EPrint Archive*, (2020/710), 1–22.

Berentsen, A., & Schär, F. (2018). A short introduction to the world of cryptocurrencies. *Federal Reserve Bank of St. Louis Review*, *100*(1), 1–16. https://doi.org/10.20955/r.2018.1-16

Biais, B., Bisière, C., Bouvard, M., & Casamatta, C. (2019). Blockchains, coordination, and forks. *AEA Papers and Proceedings, 109*, 88–92. https://doi.org/10.1257/pandp.20191018

Böhme, R., Christin, N., Edelman, B., & Moore, T. (2015). Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives, 29*(2), 213–238.

Budish, E. B. (2018). The economic limits of bitcoin and the blockchain. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3197300

de Vries, A. (2018). Bitcoin's growing energy problem. *Joule*. https://doi.org/10.1016/j.joule.2018.04.016

Deutsche Bundesbank. (2017, September). Distributed ledger technologies in payments and securities settlement: Potential and risks. *Monthly Report*, pp. 35–49. Retrieved from https://www.bundesbank.de/Redaktion/EN/Downloads/Publications/Monthly_Report_Articles/2017/2017_09_distributed.pdf?__blob=publicationFile

Easley, D., O'Hara, M., & Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics, 134*(1), 91–109. https://doi.org/10.1016/j.jfineco.2019.03.004

Eyal, I., & Sirer, E. G. (2014). Majority is not enough: Bitcoin mining is vulnerable. In N. Christin & R. Safavi-Naini (Eds.), *Financial cryptography and data security, FC 2014. Lecture Notes in Computer Science* (Vol. 8437, pp. 436–454). Springer. https://doi.org/10.1007/978-3-662-45472-5_28

Fanti, G., Kogan, L., Oh, S., Ruan, K., Viswanath, P., & Wang, G. (2019). Compounding of wealth in proof-of-stake cryptocurrencies. In I. Goldberg & T. Moore (Eds.), *Financial cryptography and data security. FC 2019. Lecture Notes in Computer Science* (Vol. 11598, pp. 42–61). Springer. https://doi.org/10.1007/978-3-030-32101-7_3

Gans, J. S., & Gandal, N. (2019). More (or less) economic limits of the blockchain. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3494434

Halaburda, H., Haeringer, G., Gans, J. S., & Gandal, N. (2021). The microeconomics of cryptocurrencies. *Journal of Economic Literature*, (forthcoming).

Hellman, M. E. (2002). An overview of public key cryptography. *IEEE Communications Magazine, 40*(5), 42–49. https://doi.org/10.1109/mcom.2002.1006971

Huberman, G., Leshno, J. D., & Moallemi, C. C. (2017). Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3025604

Irresberger, F., John, K., & Saleh, F. (2020). The public blockchain ecosystem: An empirical analysis. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3478264

John, K., Rivera, T., & Saleh, F. (2020). Economic implications of scaling blockchains : Why the consensus protocol matters. *SSRN Electronic Journal*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3750467

Ma, J., Gans, J. S., & Tourky, R. (2018). Market structure in bitcoin mining. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3103104

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Www.Bitcoin.Org. Retrieved from https://bitcoin.org/bitcoin.pdf

Naor, M., & Yung, M. (1989). Universal one-way hash functions and their cryptographic applications. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing* (pp. 33–43). https://doi.org/10.1145/73007.73011

Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies: A comprehensive introduction.* Princeton University Press.

Rauchs, M., Blandin, A., Dek, A., & Wu, Y. (2021). Cambridge bitcoin electricity consumption index. Retrieved from Cambridge Center for Alternative Finance website: https://www.cbeci.org/cbeci/comparisons

Saleh, F. (2020). Blockchain without waste: Proof-of-Stake. *The Review of Financial Studies*, (forthcoming). https://doi.org/10.1093/rfs/hhaa075

Savolainen, V., & Ruiz-Ogarrio, J. (2020). Too big to cheat: Mining pools' incentives to double spend in blockchain based cryptocurrencies. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3506748

Schär, F., & Berentsen, A. (2020). *Bitcoin, blockchain, and cryptoassets—A comprehensive introduction*. MIT Press.