New Dimensions of Information Warfare: The Economic Pillar—Fintech and Cryptocurrencies

Maurantonio Caprolu[1(✉)], Stefano Cresci[2], Simone Raponi[1], and Roberto Di Pietro[1]

[1] Division of Information and Computing Technology, College of Science and Engineering, Hamad Bin Khalifa University, Qatar Foundation, Doha, Qatar
mcaprolu@hbku.edu.qa

[2] Institute of Informatics and Telematics, National Research Council (IIT-CNR), Pisa, Italy

Abstract. The fast-paced technological advancements of the last decades have led to digitizing an ever-increasing amount of information, processes, and activities. A wide range of new digital devices have made our lives easier, faster, and funnier, quickly becoming indispensable for both work and daily life. As a result, the digital realm has dramatically expanded its boundaries, replacing the physical world in several areas. Information warfare has found fertile ground to expand into this modernized electronic world, creating new scenarios and novel attacks on nations and citizens' virtual perimeter. The economic sector plays an essential role in this context, widely affected and profoundly changed by recent technological advancements. For instance, the rapid rise of fintech systems, on the one hand, has led to the globalization of markets, with evident benefits on industries and tertiary services. On the other hand, the financial sector's dependence on digital systems and information has increased dramatically, also introducing new digital risks. This paper explores the new threats opened up by the latest technological advancements to the national economy of a typical developed Country. After identifying two of the major targets of information warfare – cryptocurrencies and stock markets – we investigate possible attacks and evaluate their potential repercussions on the national economy, also highlighting promising avenues for future research and experimentation.

Keywords: Information warfare · Fintech · Cryptocurrency · Market manipulation

1 Introduction

The end of the Cold War and the collapse of the Soviet Union in the late 80' caused tangible changes in the world economy, that engaged in progressively internationalized trades that led to the globalization of today's economy.

© Springer Nature Switzerland AG 2021
J. Garcia-Alfaro et al. (Eds.): CRiSIS 2020, LNCS 12528, pp. 3–27, 2021.
https://doi.org/10.1007/978-3-030-68887-5_1

Theoretically, this approach pursued free trade principles and a gradual disengagement of states, thereby somehow adhering to the popular theory of Montesquieu that "commerce softens manners and encourages peace" (*The Spirit of the Laws*—1748).

Sadly, history has shown us that this phase of economic globalization has quickly diverged from the principles of Montesquieu. Admittedly, as of today, free trade has essentially imposed itself, while few states have given up their political and economic supremacy prerogatives. On the contrary, economics, banking, and trade are all increasingly seen as subtle, but sharp tools of leverage and power-gathering. Indeed the economy of a nation—intended as the total production, distribution, and trade of goods and services conducted by a nation's various economic agents—is central to the livelihood of the nation. The more a nation's economy thrives, the greater the capacity of the nation to provide the public services required for the well-being of its people, including public health, education, and infrastructure, as well as military spending, which is vital to safeguarding stability against both internal and external threats.

To make an example, let us take into account the conflict between the United States of America and the Republic of China. A merciless trade war is being fought by the two nations, with major economic interests at stake. Among their weapons, it is possible to find industrial espionage, technological hacks, custom duties, and legal tools; the same arsenal that, together with soft-power, has allowed the USA to enforce their political agenda—overall, successfully, so far. A further example of a nation's interest in economic war is given by France. In 1997, France established the École de Guerre Économique (School of Economic Warfare) as an academic institution of a renowned Parisian business school, called École Supérieure Libre des Sciences Commerciales Appliquées (Free Superior School of Applied Commercial Sciences). According to such a school, the economic war is a strategy and a process decided by a state as part of the assertion of its power on the international stage, being carried out through information on the financial, economic, technological, political, societal, and legal fields[1]. In the years following its creation, the School of Economic Warfare has been proposing a curriculum based on the following assumptions: (i) the economic conflicts have been increasing during the past 20 years; and (ii) both information warfare and management are the essential means used by contestants to be predominant in such conflicts. However, given the level of complexity, both companies and nations need to boast a vast range of skills to face information warfare on the economic battlefield.

The teaching of competitive intelligence is explicitly designed to examine and resolve economic conflicts shaped by states and private companies alike. In areas such as Policy and Economic Intelligence, Risk Management, International Security, and Cybersecurity, the school currently provides postgraduate training.

[1] https://portail-ie.fr/resource/glossary/95/guerre-economique (Last checked December 2020).

One of the foundations of its model, according to the school itself was "the transfer of methodology from the military world to the civilian world."[2].

External actors, such as foreign governments and terrorist groups, may target a nation's economy in various ways and for various reasons such as undermining defensive capabilities before a military attack, or simply destabilizing a country by causing population turmoil. Indeed, it is known that a country that is destabilized and fractured is more fragile and can be affected more quickly from the outside. In this respect, any economic asset important to the nation, such as individuals, companies, organizations, or the government itself represent the attack surface. New technologies, which are constantly applied to different sectors of the economy, lead, on the one hand, to developing, optimizing, and automating economic processes, thus reducing costs and increasing income. But emerging innovations, on the other hand, eventually introduce new vulnerabilities: they raise the reach of attack and expose the economy to unprecedented risks. Consider cryptocurrencies. They have recently gained tremendous momentum and attracted hundreds of billions in capitalization.

One of the first, short contributions related to the new dimensions of Information Warfare, including the above highlighted issues, can be found in [23]. Therein, new possible scenarios are sketched, together with a coarse grain analysis of the impact of new threats on the most sensitive targets exposed by every nation: the Society, the Economy, and the Critical Infrastructures. Instead, in [24] can be found a detailed, analytical, rigorous and—to the extent possible—complete treatment of the different domains characterizing the new dimensions of Information Warfare.

1.1 Motivations

The frenetic technological progress of the last few decades is radically changing our habits and lifestyle. The subsequent digitalization of an ever-increasing amount of data is expanding the boundaries of the digital realm, exposing our society to new security challenges and risks. In this new virtual environment, cybersecurity threats can jeopardize countless new private and public assets, with potential impacts on national security hardly imaginable just a score ago. Therefore, it is not surprising that information warfare is gaining more and more strategic importance and attention from public and private industries, governments, and various other actors. Hence the need to contextualize information warfare in the current technological scenario, investigating the attack and defense techniques existing in the literature, considering different possible scenarios, and identifying open research and technology problems.

1.2 Contribution

In this paper, we delve into the novel threats introduced by the new dimensions of information warfare, specifically targeting the economic sector. We first iden-

[2] https://www.ege.fr/index.php/l-ecole/presentation/economic-warfare-school-of-paris.html (Last checked December 2020).

tified two of the most critical targets of Economic Information Warfare, i.e., the cryptocurrencies and the stock market, significantly affected by emerging security threats. We then investigated the possible attacks against these targets and highlighted the current state-of-the-art concerning existing and future threats, proposing solutions, and identifying related research and technology problems.

Roadmap. The paper is organized as follow. In Sect. 2 we present the cryptocurrencies as a target of the modern Economic Information Warfare, discussing the possible existing and future attacks against its technological pillars (Sect. 2.1) and its IT infrastructure (Sect. 2.2). We then discuss the attacks against the Stock Market in Sect. 3, investigating market manipulation techniques (Sect. 3.1), new threats introduced by the rise of high-frequency trading (Sect. 3.2), and attacks against the market's availability (Sect. 3.3). Finally, in Sect. 4, we draw some final remarks.

2 Cryptocurrencies

Blockchain-based systems, in particular permissionless ones, have a large attack surface due to the distribution, complexity, and openness of the resources involved in their protocols. The most important cryptocurrencies, such as Bitcoin, Ethereum, and Monero, are public blockchain-based systems where all users have the same permissions. Anyone can join their network, access the distributed ledger, and participate in the protocol. As a result, any user could be a potential adversary and jeopardize the security of the system. The architecture of existing cryptocurrencies requires that the system's security and consistency are verified and guaranteed by its users, without relying on trusted third parties. On the one hand, this feature allows the development of transparent systems, where each user can verify the data's consistency. On the other hand, by design, the system's security is guaranteed as long as the majority of users behave honestly. All the most important cryptocurrencies are supported by a consensus mechanism that allows the network to agree on users' transactions validity. This mechanism is based on the resources, usually computational power, that each user offers to the network to guarantee its security. To successfully compromise a cryptocurrency, a malicious user would have to own and use the majority of the total resources available in the system, performing the so-called 51% attack. Consequently, a cryptocurrency is more vulnerable in the early stages of its life cycle, when its community is still young and unstable. In fact, during the first period after its release, a cryptocurrency usually is little known and used, like any other software. In this phase, when the community's size is still limited, an attacker could easily obtain 51% of the resources and use them maliciously, compromising the system's security and consistency. This type of attack does not need to compromise resources or exploit vulnerabilities. Since the consensus mechanism is based on the majority, it is sufficient to join the network with enough resources to make decisions independently, without even violating the protocol. The 51% attack is, therefore, very effective and challenging to detect. However, it is also hard to be performed, especially against cryptocurrencies with extensive, stable,

and solid communities, such as Bitcoin and Ethereum. There are several other ways to attack a cryptocurrency, mostly exploiting the vulnerabilities of the individual modules of which they are composed. Cryptocurrencies are software systems consisting of different technologies, each of which plays a different role and allows different functions. For example, the blockchain is used to implement a distributed database that ensures data consistency through consensus among participants. In turn, peer-to-peer networks connect the nodes that make up the system among each other, enabling exchanging messages and data. In addition to their functionalities, all these technologies also introduce their vulnerabilities, increasing the attack surface. Consequently, an attacker could threaten a cryptocurrency not only by directly attacking its protocol. Malicious users could also exploit vulnerabilities or implementation errors in its software components, solve mathematical problems on which the security properties are based, and attack the underlying IT infrastructure.

The impact on a nation's economy of a possible successful attack on a cryptocurrency is highly variable. The assumption about cryptocurrency users' honesty is problematic for many people, who prefer to trust a single external entity, e.g., a bank, rather than half plus one of the other network users. The mistrust of new users towards cryptocurrencies is one of the main reasons this technology struggles to establish itself as a daily payment method, remaining much more used for investments and speculations. There are no nations that strictly depend on a cryptocurrency at the time of writing. This implies that, currently, the national security impacts of an attack against cryptocurrencies would be limited. In fact, the affected users would be companies and small investors scattered worldwide, hardly grouped in a single nation. However, several nations are dreaming about creating a state-sponsored cryptocurrency that can enhance or displace traditional fiat money. In such a scenario, the national currency would be exposed to several new cyber threats, with consequences ranging from short DoS to permanent damages to the national financial infrastructure.

In this sections, we describe some methodologies that could be used to attack cryptocurrencies, divided into two macro-categories: attacks against enabling technologies; and attacks against vulnerabilities in the underlying IT infrastructure layer. Then, we investigate how these attacks could be used to jeopardize the economy of a nation.

2.1 Vulnerabilities of the Technological Pillars

The most important cryptocurrencies, both in terms of users and capitalization, are based on blockchain. Introduced in 2008 by Satoshi Nakamoto, the blockchain is a peer-to-peer network that implements an append-only, immutable, and distributed database. The system's security is verified by its nodes, without resorting to a trusted third party.

In its original form, the blockchain relies on several technological pillars, mostly based on cryptographic functions. The Elliptic Curve Digital Signature Algorithm (ECDSA) is used, for example, to ensure that users can only spend

their own funds. Cryptographic puzzles [1], instead, are used to implement the so-called proof-of-work, a consensus mechanism that manages new block's creation and validation. Finally, storing data within an immutable chain is made possible by hash functions that concatenate the ledger's blocks. These cryptographic functions are based on particular mathematical problems, considered difficult to solve by the international scientific community, from whose complexity derives the security of the protocol. In these cases, the attacker who manages to solve the mathematical problem can break the cryptographic protocol. Among the many possible attack strategies, we can identify two very effective methodologies: reducing the complexity of the mathematical problem and using new technologies to solve it efficiently.

Complexity Reduction. The ECDSA algorithm uses the elliptic-curve discrete logarithm problem (ECDLP), a mathematical problem based on the cyclic groups of elliptic curves over finite fields, considered hard to solve. The mathematical properties used in this protocol ensure that, given a public key $pubK$, the computation of the private key pK associated with $pubK$ is infeasible. Cryptocurrencies use ECDSA to secure transactions, allowing each user, identified with a public key, to spend only their own money through the corresponding private key. This system is considered safe because deriving pK from $pubK$ is computationally too expensive for any attacker. Here, "too expensive" means that, regardless of the attacker's capabilities, the computational time spent to break the protocol exceeds the usefulness window of the violated secret. This property has been formally proved valid, as long as certain ECDSA implementation conditions are met. The sufficient conditions, identified in [8], include the following properties:

- the underlying hash function must be collision-resistance and must have the uniformity property
- pseudorandomness in the private keyspace for the ephemeral private key generator
- generic treatment of the underlying group
- a further condition on how the ephemeral public keys are mapped into the private key space.

Nevertheless, at some point, someone could either simplify the problem underlying ECDLP or create a new computational technology capable of finding a solution in a much faster time. In both cases, the opponent would be able to calculate the other users' private keys, thus becoming able to spend their money. Victims would have no opportunity to get back the stolen crypto money, as transactions are irreversible in the blockchain environment. The possibility of anyone reducing the complexity of ECDLP is considered very unlikely. However, it would not be the first time that a mathematical problem, considered difficult for many years, has been solved. Fermat's conjecture, formally known as Fermat's Last Theorem, is a famous example of such an eventuality. This theorem, formulated in 1637, asserts that the equation $x^n + y^n = z^n$ does not admit solutions for

integers $n \geq 3$. Although plausibly correct, this conjecture remained unproven for three centuries when in 1994, a British academic, Andrew Wiles, published a formal proof. Just as happened with this conjecture, advances in mathematics or technology could involve solving problems underlying Elliptic-curve cryptography (ECC) or ECDSA, exposing new vulnerabilities able to break the security of current cryptographic protocols.

Other examples of similar eventualities are the "baby-step, giant-step" algorithm, and Pollard's rho method. Although not aiming to decrease the mathematical complexity of the problem, these two algorithms have tried to solve ECDLP using "shortcuts" compared to classical solutions. Nevertheless, although significantly optimizing the resolution of the problem, these algorithms do not yet allow to attack ECDS in a reasonable time, i.e., fast enough to threaten the security of the systems that use this cryptographic protocol. Unfortunately, there is no way to predict if and when further optimization of these algorithms could be discovered and used to compromise cryptographic protocols, leaving tremendous uncertainty about these technologies' future security.

New Technologies. As discussed above, existing cryptocurrencies rely on cryptographic protocols to guarantee the security of the network. These protocols are proven safe against any adversary, regardless of their abilities. However, an attacker with unexpected computational power, not available at the time of cryptographic protocols' design, may be able to take an unfair advantage over other users. An example of such a scenario concerns cryptocurrency mining and the introduction of ad-hoc hardware: the Application-specific integrated circuit (ASIC). Before the advent of ASIC hardware, the computational power made available by users to secure the Bitcoin network came only from generic-purpose hardware. No user had a consistent advantage over the others. With the release of ASIC hardware, specifically designed to optimize mining activities, the Bitcoin network balances have changed. Users who started mining with ASIC have had such a massive advantage in computational power that mining activities with generic hardware has become ineffective and unprofitable. Since this hardware was created to be immediately distributed on the global market, the beneficiaries of this novel technology were numerous. Consequently, the new computational power is widely distributed for users and geographic areas, as depicted in Fig. 1, avoiding its centralization on a single entity.

Conversely, suppose this technology was not intended for the global market. In that case, its developer could have used it to gain a computational advantage over other users, jeopardizing the network's security. Therefore, the large-scale distribution of new computational technologies is essential to avoid problems of stability and security of cryptocurrencies. However, this may not always be possible. In the case of Quantum Computing (QC), for example, the high production costs could slow or prevent its distribution on the global market. As a result, the manufacturing company could be the only one in possession of such computational power, gaining such technological supremacy as to allow it to control current cryptocurrencies. At the time of writing, we are still a long way from

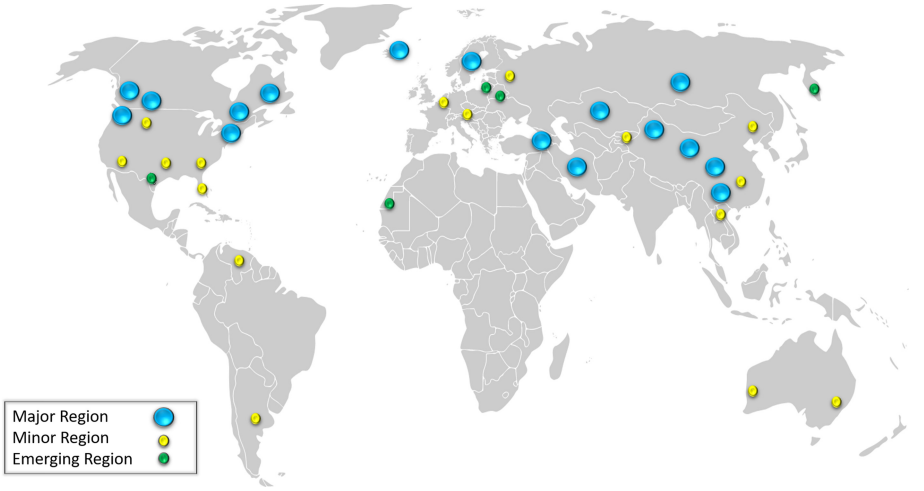

Fig. 1. Global overview of the Bitcoin mining regions. Data sourced from [6]

obtaining a quantum computer capable of endangering current cryptographic protocols. However, research is proceeding rapidly in this area, reaching increasingly important milestones. According to initial results, the incredible computational capabilities of quantum computers promise to perform tasks, infeasible in today's computers, efficiently. An attacker could use this technology to break current cryptographic protocols, seriously endangering systems, such as cryptocurrencies, which base their security on these mechanisms.

2.2 Vulnerabilities of the IT Infrastructure

Another type of attack against cryptocurrencies aims to exploit vulnerabilities in the underlying IT infrastructure, such as software modules or network infrastructure. Vulnerabilities in software modules, such as electronic wallets, blockchain management software, or the transaction validation system, can be exploited to harm individual users or the entire network. The vulnerabilities of the network infrastructure, on the other hand, can be exploited to compromise multiple functions, such as the consensus mechanism, by tampering with messages in order to alter interactions between users.

Software Vulnerabilities. Although the mathematical properties on which ECDSA bases its security have never been compromised, we can cite several examples of attacks that exploited vulnerabilities in its implementation. One of the most sensitive phases of implementing a cryptographic protocol is the choice of the security parameters. In ECDSA, for example, the choice of the elliptic curve and its domain parameters determines the robustness of the encryption keys produced. The scientific community has extensively studied different types

of elliptic curves, releasing standard parameters universally accepted as safe when correctly used in the implementation of ECDSA. However, each developer is free not to use the standard parameters, replacing them with other customized versions. Besides, the scientific community has also studied the best practices to be followed during the implementation of ECDSA and other cryptographic algorithms. Some particular parameters, for example, must be chosen randomly at each execution of the protocol. Their static setting could introduce severe weaknesses in the generated encrypted material, as happened to Sony in generating their key pair used to digitally sign video games for their Playstation3 console [35]. They used a static parameter (rather than random) to implement the ECDSA protocol, making the resulting private key computable by analyzing few digitally signed files. This flaw was discovered by a group of hackers, known as *fail0verflow*, who was able to reconstruct Sony's private key and use it to distribute counterfeit video games.

Often, the software implementation of a cryptographic protocol could be vulnerable even if developers diligently follow the best practices suggested by the scientific community. Usually, one of the most significant problems is how to generate random numbers, especially on mobile device platforms, where the available resources are limited. There are several libraries capable of generating pseudo-random numbers. [41] provides an investigation on the most used Java libraries to generate random numbers, evaluating the methodologies used and the quality of the numbers generated. They found multiple flaws on entropy collector components, with different severity and probability of occurrence. In details, they showed that the Android PRNG's overall entropy could be reduced to only 64 bits. This flaw was exploited in 2013 to steal Bitcoin from accounts generated by electronic wallets in the android environment. The Java class SecureRandom, used by unsafe digital wallets, has been identified as the main responsible for the introduced vulnerability, generating collisions on the produced random numbers. The best practices require that the random number used to sign a private key in the ECDSA protocol cannot be reused. If the randomly generated number is used more than once, the private key could be easily computed by an adversary.

Another severe vulnerability discovered in 2011 enables a full key recovery attack against a TLS server that manages authentications with ECDSA signatures. As described in [9], a vulnerability in OpenSSL's implementation allows a timing attack affecting the generation of the encryption keys used for the digital signature.

Network Hijacking. In this type of attack, an adversary maliciously interferes with a cryptocurrency protocol manipulating the network traffic used by honest users to communicate. The attacker can be either an external or an internal user. In the first case, the attacker does not participate directly in the protocol but performs a passive attack. A classic example would be an attacker dropping network packets, preventing specific users from exchanging information with the rest of the network. In the case of an insider, however, the attacker joins the network and participates in the protocol as an honest user, interacting publicly

with other nodes. Then, he begins to behave maliciously, sending false information, creating fake transactions, or acting in any other way that does not comply with the protocol.

Several factors influence attacks on the network infrastructure of a cryptocurrency, conditioned by the technologies used by the single protocol. Therefore, cryptocurrencies are not all affected in the same way by this type of threat. Furthermore, the vulnerabilities are often due to external conditions, not managed by the protocol or the technologies used by cryptocurrency, such as the Internet routing infrastructure.

By design, every permissionless cryptocurrency allows anyone to access the network without authentication. To participate in the protocol is sufficient to run a full node, from anywhere in the world, with an active Internet connection. For this reason, how ISPs manage their network directly affects cryptocurrency full nodes' ability to communicate with each other. It is also important to note that, although anyone in the world can join the network, it is very unlikely that the nodes are geographically distributed in a uniform way. As a direct consequence of this, full nodes are likely to be grouped in a few regions, hosted under the infrastructure of a few ISPs, that will be responsible for routing the entire network's traffic. In this scenario, multiple attacks may be carried out, either actively or passively, to threaten the cryptocurrency by targeting the ISP's network infrastructure. We list in the following a few malicious behaviors that may be executed either by an external attacker or by a malicious ISP:

- *Network traffic redirection:* by exploiting vulnerabilities in network protocols, i.e., bgp hijacking.
- *Network traffic filtering:* by maliciously dropping selected packets, causing DoS, i.e., Blackhole attack.
- *Network traffic manipulation:* to isolate specific nodes only, i.e., Eclipse attack.

These malicious operations, known as Internet routing attacks, could be used, alone or together, to carry out the following attacks:

- *Partition attack:* The goal is to split the peer-to-peer network of the targeted cryptocurrency into separate disjointed segments, such that the different sections are no longer able to communicate.
- *Delay attack:* The goal is to postpone the spread of new blocks through the network to enable multiple other attacks, such as double-spending.

All of the attacks listed above can be performed against any cryptocurrency that relies on public internet infrastructure to manage inter-node communications. The motivations behind these attacks, as well as the consequences, could be manifold, while the possible impacts vary according to the victim. For example, a storekeeper could be subject to temporary outbreaks that prevent his activities, as well as more severe issues such as the double-spending attack. If under attack, miners could waste the computational power that they provide to the network to guarantee its safety, facing lost earnings. Finally, a regular user would face DoS attacks that prevent the access and use of the payment services.

3 Stock Market

The majority of existing studies at the intersection of security and economics focused on problems of micro-security – that is, how to enforce security for specific applications and protocols, or how to protect data about users of a given service. This approach is orthogonal to that focusing on macro-security, which concerns with the security and trustworthiness of whole markets and technologies. While the former is of great importance and concern for individual users, the latter is instead primarily of interest for governmental actors and nations themselves, thus falling under the broader information warfare umbrella. This is due to the potentially significant influence that macro-security threats can exert on the national economies. Within the context of economic war, nations regard the economy as a worldwide arena and the latest technological advancements as sharp weapons with which to advance their strategic and political agendas. Here, the vulnerabilities inevitably introduced by such new technologies, such as those that contribute to the rise of fintech, combined with the weak regulatory frameworks, can be profitably exploited by malicious actors. The existence of many ways to directly compromise fintech services, or to tamper with their underlying technologies, means that an opponent nation could use the very same fintech technology as a weapon. Indeed, fintech services can easily become attack vectors that could lead to the compromise of critical economic resources of competing nations. In this regard, we can easily find a plethora of news on alleged state-backed actors and state-sponsored hacking on newspapers and information sites [21].

Among the paramount examples of the systems and technologies that contribute to the rise of fintech, are the national stock markets. Stock markets, or equity markets, are one of the most important economic assets of a nation and a constituent of national free-market economies. They refer to centralized physical or virtual spaces where equities or stocks of publicly held companies, bonds, and other classes of securities, are issued and traded. Given the central and crucial role that stock markets have within the economic processes of a nation, fair and secure operations should be guaranteed at all times. However, while in the early days of physical hectic trading floors this posed somewhat manageable challenges, the security risks introduced by the wide array of technologies that permeate current physical and virtual stock markets have escalated to new – dangerous – heights. Among them, are the risks related to the many different existing forms of market manipulation, aimed at artificially inflating or deflating the value of given traded securities. When targeted at country-relevant stocks or nation-critical firms, these forms of market manipulation are capable of endangering a whole national economy. Then, new rapidly-evolving technologies such as automatic trading (AT) and high-frequency trading (HFT) are responsible for a progressively larger share of market transactions. Their role within catastrophic flash crashes and their effects on market stability are still debated, so as their contribution to stock market security concerns. Finally, the recent COVID-19 pandemic sped up the ongoing virtualization and "remotization" of trading floors and stock markets. With always less manual intervention in favor

of remote software-mediated operations, a new wave of security threats needs to be addressed.

In the remainder of this section, we investigate the different ways in which fintech and the stock market can be weaponized to attack a nation's economic assets, describing the current state-of-the-art with regards to both attacking and defensive means.

3.1 Market Manipulation Threat

The new market manipulation methodologies share the same objectives as their traditional techniques. However, efficiency has improved dramatically as new manipulation methodologies leverage the latest technological advances and operate in a different, faster, and highly interconnected digital market [57]. There are different forms of market manipulation. Some of these aim at marginal, low-value stocks, while the more aggressive aim to hit the heart of the financial market. These latest forms of manipulation have the potential to create massive shocks in national and global markets, making such activities a primary national security concern for any country.

Previous studies on this subject classified manipulations into two main categories: (i) information-based, and (ii) trade-based. Information-based manipulation consists of distributing false information or publishing fake news to have a specific effect on the stock markets. On the contrary, trade-based manipulation is based solely on shares' movement, without involving other publicly observable information such as disseminating fake news [57]. Some of these manipulation techniques have always existed. However, in recent years, they have become increasingly widespread, effective, and indistinguishable from legitimate actions, thanks to the recent technological progress.

Information-Based Manipulation. Traditionally, stock market forecasts were obtained by exploiting historical stock market data, for instance, by training statistical autoregressive models. In recent times, however, it has become clear that also other types of information could be used to predict future market trends. A new quest thus began to discover and exploit other informative data sources. Among them, textual data from official news outlets and spontaneous user posts in online social platforms showed great potential and predictive power [44]. For instance, news articles from Yahoo Finance are leveraged by the system proposed in [50] to predict future prices of S&P 500 stocks. As another example, asset volatility movements are predicted in [2] by processing information extracted from several news sources. The work in [25] applied text-based event detection to identify noteworthy events. Such events are then fed to deep convolutional neural networks to model both short-term and long-term influences of events on stock price movements. For what concerns data extracted from online social networks, it has been shown that the sentiment polarity of user posts holds great predictive power for forecasting movements in financial markets [7,54]. The system described in [11] is a notable example of this body of

work, where authors trained a machine learning classifier capable of predicting the next day trend of certain stocks by only exploiting the sentiment value of stock-related tweets. In a similar fashion, opening and closing prices are predicted in [49] by leveraging sentiment analysis of social media posts. In addition to the mere textual data, also other types of data extracted from online social networks can be profitably used for market prediction. For instance, in [34] authors developed methods for market prediction and for portfolio selection by leveraging correlations between companies that co-occur in social media posts. Co-occurring companies are modeled via large networks of companies and results are obtained by the application of graph mining techniques.

All the previous examples highlight the growing importance of alternative online data sources for stock market prediction. However, serious concerns arise if we consider the possibility and the ease with which online data can be tampered with, manipulated and even outright fabricated [17]. Thus, on the one hand, many systems for monitoring and predicting stock markets are now heavily based on the analysis of online data. On the other hand, however, a significant share of such online data turns out to be fake, inaccurate and misleading, thus possibly leading such systems astray. Should this risk materialize, consequences in terms of market crashes and widespread financial losses would be dramatic, as it already happened in a few notable cases [27]. In recent years, this risk motivated an emerging stream of research on online financial disinformation, which already led to interesting – yet worrying – findings. Among the most striking results, is the detection and investigation of an online manipulation campaign carried out on the Twitter microblogging platform [16,17]. In detail, the authors analyzed some 9 million tweets mentioning 30,032 different companies traded in the main US financial markets (e.g., NASDAQ, NYSE, NYSEARCA, NYSEMKT, OTCMKTS). Within this dataset, they found suspicious co-occurrences between a few stocks with very high market capitalization and many unpopular stocks with very low capitalization. By applying state-of-the-art bot detection techniques [15], the study found that more than 70% of all users that tweeted about the low-capitalized stocks, were in fact bots – namely, automated accounts used for large-scale spamming [14]. Going forward, a subsequent study also analyzed the characteristics of such financial bots, concluding that their goal was that of luring automatic trading algorithms into buying the low-value stocks by exploiting the popularity of high-value ones [52]. The need for evaluating the credibility of stock-related social media messages is also underlined in [26]. These findings currently represent the first large-scale, empirical evidence of widespread financial spam in online social networks.

Currently, no system exists that is specifically designed for detecting online financial spam and financial manipulation campaigns. In fact, the more generic problem of detecting online information manipulation is already very challenging, with few existing solutions that demonstrated decent performance. As such, at the time of writing, protection against online financial disinformation must necessarily rely on techniques for defending against generic information manipulation. Successful disinformation campaigns are those that manage to reach and

influence a large number of users. To achieve this goal, perpetrators typically leverage large numbers of automated (i.e., bot) or paid (i.e., troll) accounts, in order to reshare and broadcast their malicious messages [14]. Based on this consideration, a first line of defense against information-based market manipulation revolves around the application of bot and troll detection techniques. A recent survey on the topic highlighted that, among the plethora of existing approaches for detecting malicious accounts, those that are based on unsupervised approaches for the analysis of suspicious behavioral similarities are the ones that manage to obtain the best detection performance [14]. Examples of this kind are [12,15,33,38,39]. This is in contrast to earlier approaches based on the application of supervised classifiers that analyze each account individually. The survey in [14] also underlined the importance of accounting for adversaries, motivated in evading detection systems, by design. This can be obtained by designing detection systems that leverage recent advances in adversarial machine learning, such as in [55], or anyway by adopting adversarial approaches to the study and detection of malicious accounts, as done in [18,19]. Another emerging and promising direction of research is aiming to detect so-called coordinated inauthentic behavior (CIB). Also, in this line of research, the focus is posed on coordinated accounts. However, when studying CIB the nature of individual accounts (e.g., whether they are human- or software-operated accounts, bots, trolls, etc.) is not of interest anymore and the only dimensions that are deemed meaningful are coordination and authenticity of the online personas and of the content they share [46,48]. Finally, yet other approaches to contrast online information manipulation are related to the computational detection of fake news and propaganda at scale [20,58].

Trade-Based Manipulation. In contrast to information-based manipulation, trade-based manipulation attempts are exclusively based on buying and selling shares, without requiring to share false or misleading information. These types of market manipulation are as old as the markets themselves. However, they recently regained widespread attention as a consequence of the rise of new technologies. In fact, while the vast majority of long-established stock markets enforce strict regulations for avoiding trade-based manipulations, some of the newest and less regulated exchanges provide fertile ground for such nefarious practices to proliferate once again. Among the newest and less regulated exchanges are cryptocurrency exchanges [47]. In addition to the widespread adoption of, and demand for, cryptocurrencies, also other relatively new technologies facilitated the furious comeback of trade-based market manipulations. As in the case of disinformation campaigns, also trade-based manipulations necessitate large numbers of (aware or unaware) participants to be involved, in order to achieve substantial results. As such, online social networking platforms – characterized by the sheer number of users and by the large support for anonymity – again represent a profitable avenue for manipulators. This is the reason why several scholars recently devoted significant efforts towards the study of online cryptocurrency manipulations [47].

Among the most widespread and potentially detrimental trade-based market frauds, are pump-and-dump schemes and Ponzi schemes. Specifically, pump-and-dump involves the artificial inflation of the price of an owned stock, with the goal of selling it at a higher price. The perpetrators of this fraud typically buy low-value coins way before the scheme takes place. Then, they lure other willing participants and unaware investors into buying the stock, thus causing a surge in price. In turn, this surge inevitably attracts other investors thus raising the price even more. When the price reaches a given target value, the initial participants simultaneously sell. Shortly after, the other aware participants sell as well, thus starting a price collapse. In a matter of minutes the prices plummets, reaching values that are way lower than the initial ones. As a result, a few organizers and early participants manage to obtain large gains, while the vast majority of other aware and unaware investors suffer severe losses. When planning pump-and-dump schemes, orchestrators typically target small, thinly-traded coins, since it is easier to manipulate princes when there is little or no independent information available about the security, or little activity anyway. Based on the above description, attracting many investors is instrumental for successfully orchestrating a pump-and-dump scheme. Traditionally, unaware participants were lured by using spam e-mails, fake press releases and via telemarketing from "boiler room" brokerage houses. However, in more recent times, online financial discussion boards, social networks and messaging apps are the media of choice for attracting participants.

Given the pivotal role of social media in trade-based frauds, a growing number of studies focus on characterizing online social media discussions about cryptocurrencies with the aim to uncover possible manipulations. Among such studies is [29], which investigated Reddit discussions about the Bitcoin, Ethereum, and Monero coins. Authors found that Monero, in particular, is often used for shady or illicit activities. Interestingly, they also measured longer and wider information cascades for Monero, with respect to those of the other coins, showing that many of the users interested in cryptocurrencies are actually interested in coin manipulations [29]. Instead of focusing on a specific platform or set of coins, the work in [47] adopts the first large-scale and context-agnostic approach to investigate online cryptocurrency manipulations. In detail, authors collected a large multi-platform dataset including conversations about a multitude of coins from Twitter, Telegram and Discord, including both genuine cryptocurrency discussions as well as those about cryptocurrency frauds. The large-scale analysis managed to uncover a large number of previously unknown Telegram channels and groups, some of which were invite-only, specifically dedicated to organizing and coordinating pump-and-dump schemes. Contrarily, Discord appeared as a relatively safe platform, for what concerns cryptocurrencies. The study in [47] also allowed to detect hundreds of automated Twitter accounts that are used for advertising ongoing pump-and-dump operations, thus luring unaware investors into the fraud. In addition to the previous studies, a few other works specifically focused on an online platform or coin. As an example, the work in [56] focused on Telegram pump-and-dump schemes, by providing a detailed account of how

such frauds unfold on the platform. As a by-product of their analysis, the authors also developed a simple machine learning classifier for predicting the likelihood of a coin being the target for manipulation. Then, they tested a simple trading strategy that invests in coins with a high likelihood of being pumped in the near future. Notably, their results showed that the simple trading strategy allowed to obtain a return as high as 60% over the course of two and a half months. Adding to the observational studies previously summarized, the work in [42] proposed a first inferential analysis. In particular, authors built and leveraged a Telegram pump-and-dump dataset to train machine learning models for solving a number of tasks. The first task that they experimented with aimed at detecting pump-and-dump scams as they unfold, based on the sequence of messages shared in a Telegram group/channel. Then, they also developed a model for estimating the likelihood of a given pump-and-dump attempt to succeed. Within this context, a pump-and-dump is considered successful if the pumped coin manages to reach the target price set by the organizers. As a final result, they also investigated the presence and role of Twitter bots in cryptocurrency-related discussions. Their results confirm earlier findings in [47], showing a large prevalence of bots during coin pumping operations.

Among the other forms of trade-based market manipulation, also Ponzi schemes received scholarly attention. Ponzi schemes – named after the infamous swindler that orchestrated the first of these scams – are investment plans that promise extremely high rates of return in very limited time. In reality, however, participants' money is not invested, but instead it is used to provide returns for earlier backers. Similar to other pyramid investment schemes, in order to be sustainable, also this scheme necessitates a constantly growing inflow of money, to be obtained from an equally growing number of participants. As such, Ponzi schemes eventually bottom out and unravel when the flow of new investors isn't enough to sustain the scam. Based on the aforementioned functioning of this manipulation, perpetrators typically devote all of their efforts to attracting new participants. As in the case of pump-and-dump schemes, social media allow scaling the recruitment of new participants to Ponzi schemes to a whole new level. The study in [47] investigated the presence of Ponzi schemes in Twitter, Telegram and Discord discussions. Authors found no evidence of users involved in Ponzi schemes on Discord. However, they found tens of Telegram channels and groups specifically devoted to these scams. A peculiarity of these channels and groups was that they all pointed towards one another. In fact, every of such channels contained links for joining channels and groups related to other Ponzi schemes. While, on the one hand, this makes it easier for the orchestrators to recruit new investors via mutual advertisement, it also allows to trace back and identify the majority of channels involved in this manipulative practice, with relative ease [47]. Ponzi schemes have also been investigated within the Bitcointalk online discussion board [53]. Authors leveraged techniques for survival analysis with the goal of identifying the key factors that determine the success of Ponzi schemes. Others also proposed a machine learning classifier for detecting such schemes, by analyzing features derived from the Bitcoin blockchain [4].

3.2 Double-Edged Technologies

One fundamental dimension of technological advancement in fintech is undoubtedly represented by *speed*. Today, market transactions are issued and resolved in a matter of microseconds and at unprecedented volumes, thanks to High-Frequency Trading (HFT) and Automatic Trading (AT). AT identifies trading systems that leverage software algorithms to automatically determine orders to issue, modify or withdraw, with limited or no human intervention. HFT represents a specialization of AT that also introduces dedicated infrastructures to minimize network and computation latencies by leveraging specific facilities such as co-location, proximity hosting, high-speed direct electronic access and high-performance computing [43]. By exploiting these advanced technological means, HFT is capable of monitoring prices and transactions across many different global markets at the same time. In addition, it also allows to establish and liquidate positions in very short time-frames, based on such aforementioned real-time market conditions. HFT is thus regarded as an advanced technology capable of opening up new trading possibilities for its adopters, by benefiting from lightning-fast analyses and transactions with respect to the slower traditional traders. These unprecedented capabilities result in the possibility to take advantage even from minor price differences. As a result, high-frequency traders frequently benefit more from a large number of minor transactions than from a few particularly significant ones, as manual traders do [43].

The Role of HFT Under "Stable" and "Critical" Market Conditions. In light of the disruptive changes introduced by HFT and its sometimes shady uses (e.g., arbitrage, front-running), a large body of work investigated the role and effects of AT and HFT on stock markets. One notable finding emerging from the analysis of the existing literature is that the vast majority of existing studies reported overall positive market effects for the adoption of HFT in stable markets (i.e., when markets are not undergoing a crisis or crash). In detail, it has been documented that HFT contributes to the reduction of information asymmetry between buyers and sellers. Some studies also empirically verified that HFT contributes to market liquidity and to shrink intraday price volatility [21]. Other examples also provided evidence that HFT may contribute to stabilize markets [32], to improve market quality by reducing the bid-ask gap [31], and to reduce trading costs [40]. In summary, all these results hint at the possibility that HFT plays a relevant beneficial and stabilizing role for markets, when these operate in stable conditions. In turn, this finding suggests that regulatory measures designed for hampering the activities of high-frequency traders could in fact lead to negative market consequences, especially in terms of market liquidity [43].

The previous positive results are all related to the adoption of HFT in markets that operate under "normal" conditions. However, opposite results were obtained when analyzing markets during distressed times, as for example in the case of flash crashes. Starting from the infamous 2010 Flash Crash, several studies documented a negative role of HFT in initiating and amplifying market crashes [13]. To this end, some authors found evidence that HFT tends to

exacerbate transient price impacts that are unrelated to fundamentals – a situation that is typically observed during a flash market crash [5]. The key message emerging from the still-growing body of work that examined the role of HFT in distressed markets, is that it acts as a catalyst for existing market dynamics, including bubbles and crashes. The growing interdependencies between disparate financial instruments are likely to lead to even more frequent and complex market crashes in the future. In this rapidly evolving scenario, the technological arms race that is peculiar of AT and HFT could favor the emergence of catastrophic market crashes [51].

Technological Bias and Monopoly. In the previous sections, we highlighted the role played by HFT in generating market crashes. We also addressed the powerful connection between HFT and its underlying technologies, which determines its unprecedented speed and performance. Worryingly, the combination of HFT technology and flash crash opens up new scenarios that give state actors the possibility to carry out market manipulation to strengthen their economy or weaken enemy nations. If the best performing and faster technologies are widely available and almost evenly distributed across all actors in a financial market, no single agent could hold a significant advantage over the others. Nevertheless, a specific entity could obtain a substantial and illegitimate advantage if it succeeds in developing or acquiring a much more efficient technology than those owned by the opponents. The main open problem regarding the possible weaponization of HFT for information warfare is thus related to technological bias and monopoly. Technological bias can be defined as the asymmetry or imbalance in the technology that is available to different economic actors operating within a system. The technology level has never been perfectly balanced between the various players in the stock market. However, if the technological capabilities are too unbalanced, the repercussions on the financial markets can quickly become critical. If the technological asymmetry widens to the point of leading to a technical monopoly, the involved entity could even find itself able to lead the market.

To the best of our knowledge, up to now, nobody has exploited the technological bias in HFT to put on attacks against the national economies and assets. Furthermore, although despite the importance it holds and continues to gain, technological bias failed to attract the interest of the academic world. However, it managed to draw attention from other stakeholders, often directly exposed to the dynamics of the market, including market traders and the state decision-makers. To give an example, the so-called "slow traders" have been avoiding markets that are polluted by high-frequency traders, since they would be overwhelmed. To help slow traders to avoid HFT, numerous finance professionals are continuously debating about changing the structure of the market. As a result, some famous firms are currently basing their business on providing this kind of information, for instance, by developing big data and deep learning platforms that provide daily estimates of aggressive high-frequency traders across different markets.

In addition to HFT, other areas of fintech reported the negative effects of technological bias. For example, we already covered the noteworthy case of ASIC hardware for cryptocurrency mining in Sect. 2.1. In addition, also the improvements that Artificial Intelligence and, more specifically, Deep Learning are bringing to the market forecasting are often considered as another potential factor for technological bias. The application of these powerful techniques may also create several challenges for the efficiency of the market, together with information asymmetry and irrationality of decision-making. The technological division that is thus taking shape can be leveraged by skilled traders for netting excess returns, at the expense of traders who are used to adopting more traditional technologies [28]. In the same paper, the author reports the results for Forex tradings, in contrast with the efficient-market hypothesis. According to the study, the progressive enhancement in computational software and methods will improve the trading strategies of the individual, with the obvious consequence that some traders will be more successful than others, contradicting the classical definition of a market with perfect competition. Nonetheless, it adheres to the adaptive-market hypothesis [37] that sees the market as fiercely competitive ecosystems rather than efficient ones. Given the changes the market undergoes over time, numerous adaptation mistakes can occur, mostly consequence of the different degrees of adaptation of the participants. As a result, more significant returns are obtained by some of them when compared with the others. In this scenario, technological innovation represents a primary driver for change in the ecology of the market [28].

The considerations above apply for direct harms – e.g., immediate financial loss due to both automatic and high-frequency trading, but indirect consequences, for example the diminished confidence in financial markets, are also raising a lot of attention, potentially having a bigger (and worse) impact. Other than changing those who can be harmed by trading, high-frequency trading changed how they might be harmed, and the scale of the harm [22]. Accordingly, the loss of confidence derived from failures and systemic crashes may curtail the investors' appetite for risk, thus resulting in slower (or worse, stalling) economic growth [22]. To support this hypothesis, the authors took into account the Knight Capital Group case. The firm lost $440 million in less than 30 min on August 1st, 2012, because of its new automatic trading software. This software flooded the market with orders thus forcing the temporary closing of the New York Stock Exchange. The harm caused to both the firm itself and its shareholders was tragic and almost led to bankruptcy, other than having a huge indirect impact on both the investing public's confidence and in the structure of financial markets.

Possible countermeasures to the previous issues are still under discussion, and existing proposals are coming primarily from the regulatory and ethics communities. Both computer scientists and engineers seem not to work on possible countermeasures, thus motivating the fact that technical papers discussing security issues of high-frequency trading are lacking. Taking into account regulations, some of the proposed solutions have the goal of de-powering high-frequency trading by

changing the way markets evade pending orders. Some have argued that the priority rules determining the sequence of execution of the orders that have been submitted are designed to give priority to speed. However, the regulatory conundrum is whether the time-price priority disproportionately rewards high-frequency traders and leads to risky over-investments in the technology arms race [3]. The main benefit of the currently adopted priority rules is the fair treatment of every order. Nonetheless, other priority rules have been proposed. To make an example, a rule allows every order at a price to get a partial execution, regardless of the time [36]. Others have proposed to replace the continuous trading model with periodic auctions, which can be devised to both minimize the speed advantage and mitigate other negative outcomes coming from continuous trading (e.g., manipulative strategies) [10]. As the primary benefit, the adoption of periodic auctions would allow to reduce the trading speed and to eliminate the arms race for speed. Several markets may already boast auctions at the open and close times, and are considering the introduction of midday auctions, besides the continuous trading segment [36].

Apart from the previous countermeasures, some politicians hinted at the opportunity to introduce other initiatives. To make an example, Hillary Clinton suggested introducing a small tax on the cancellation orders, with the aim of trying to crush the practice of spoofing[3]. The introduction of taxes to financial transactions, however, would face enormous difficulties, also due to the undesirable consequences and the potential risks that such an action may cause [30]. Conversely, specific taxes aiming at thwarting high-frequency trading are seen as a more sensible and desirable possibility, although being difficult to implement [36].

3.3 Threatening Availability

Stock markets prove to be determining players in the modern economy panorama, allowing easy accesses and allocations of capital to the citizens and supporting the stabilization of security prices. A multitude of financial services is offered by stock markets, which can be seen as their hub. For this reason, denying or even only limiting access to these services may have dreadful impacts on the national economy. Even individual citizens, in case of interruption of the service, are immediately affected. An example is given by the widespread panic reaction caused by the suggestion of the possibility of a market holiday in the United States, as well as by the suspended trading in other countries. As with cryptocurrencies, the physical to the virtual transition of the stock market is also critical and introduces a series of security challenges that need to be addressed. Being the stock market fully-online, the first concern that comes into mind is related to its availability.

Denial of Service (DoS) attacks are among the most common types of cyber-attacks that aim at limiting the availability of a resource to users. These attacks

[3] https://www.cnbc.com/2016/07/22/hillary-clintons-financial-transaction-tax-why-it-may-not-work.html.

are carried out by malicious actors by overwhelming the target resource with fictitious requests, thus preventing some (or worse, all) legitimate requests to be satisfied. When the Denial of Service attack is carried out in a distributed fashion (i.e., the incoming traffic flooding the victim is originated by many sources), it takes the name of Distributed Denial of Service (DDoS). With respect to DoS attacks, DDoS attacks are more difficult to defend against, since there is a multitude of machines to defend against, rather than a single one.

Denial of Service. Denial of Service, as well as Distributed Denial of Service attacks, are usually perpetrated for profit (i.e., ransom to get the service availability back), for obtaining advantage on a competitor, or for ideological reasons. However, there have been cases in which state actors are involved in DDoS attacks for both political and economic reasons. An example is given by the DDoS attack on Estonia in 2007, targeting government services, financial institutions, and media outlets. The impact was devastating, since Estonia was an early adopter of e-government and was almost paperless at the time, enough to have needed to hold the national elections online. For many, this attack is considered to be the first case of cyber warfare in response to the political conflicts between Russia and Estonia, with the former suspected to be the perpetrator[4]. A more recent example involves the 2019 Hong Kong protests against China. During the conflict, the notorious instant messaging app Telegram suffered a large scale DDoS attack, with the aim of preventing protesters from coordinating their efforts. Detailed investigations by Telegram made it possible to understand the origin of the attack, that seems to be carried out by a State-sized actor via IP addresses originating from China[5].

The aforementioned examples show how state actors have the opportunity to weaponize cyber attacks with the aim of satisfying their economic and political goals. Although, until now, no records of state-driven attacks against national stock markets exist, partly due to their physical component, with the gradual dehumanization of stock markets this scenario might promptly change. Considering the sensitivity of the markets to uncertainty, the trading interruption, even for a limited period of time, could cause a sharp fall in stock prices. The online components of stock markets, as well as the ones of other financial institutions (e.g., online banks), among other things, are not new to attacks aimed at undermining their availability, carried out both by hackers and fraudsters.

In 2013, the International Organization of Securities Commissions (IOSCO) published a report with a survey of 46 stock exchanges [45], detailing that more than half of them had already been victims of Denial of Service cyberattacks that year. Most of the attacks considered did not have effects on the functioning of the market itself and caused only less than $1 million costs for the targeted market. A couple of attacks that are worth mentioning are the one against the NASDAQ, NYSE, and BATS stock exchanges in the United States and the one against the

[4] https://www.cloudflare.com/learning/ddos/famous-ddos-attacks/.
[5] https://www.pcmag.com/news/chinese-ddos-attack-hits-telegram-during-hong-kong-protests.

Hong Kong Stock Exchange, which overwhelmed its website and heavily affected its ability to both publish filings and display prices. Furthermore, an attacker may have the opportunity to preemptively buy (or sell) shares on a market with the aim of increasing (lowering) the value of the manipulated shares, thus obtaining an immediate biased profit from its move. This is possible to achieve by either targeting a specific company, thus shaping the price of its stocks, or by targeting a specific market, thus causing a flash crash, with potentially nefarious repercussions on whole national stock markets.

4 Conclusion

Economy is among the most important dimensions affected by information warfare, since nations and other state-actors are increasingly interested in exploiting economic leverages to pursue their strategic goals. In this work, we discussed and surveyed the scientific frontier of economic information warfare, specifically focusing on two fundamental technologies – cryptocurrencies and stock markets – that are particularly affected by emerging security threats. Each of the cited topic currently represents a salient along the vast scientific frontier of economic information warfare. For each technology, we highlighted the current state-of-the-art concerning existing and future attacks as well as the possible countermeasures to contrast them. In detail, we discussed threats to cryptocurrencies both with respect to their mathematical and technological foundations (e.g., attempts at breaking elliptic-curve cryptography) as well as their underlying IT infrastructure (e.g., software vulnerabilities and network hijacking). For what concerns stock markets, we discussed the main tools for market manipulation, either information- or trade-based. In addition, we also investigated the new threats introduced by the rise of high-frequency trading (HFT) and by remote stock markets. Finally, we also highlighted some promising directions that can contribute to safeguarding our critical economic systems from the growing threats of information warfare.

References

1. Ali, I.M., Caprolu, M., Di Pietro, R.: Foundations, properties, and security applications of puzzles: a survey. ACM Comput. Surv. (CSUR) **53**(4), 1–38 (2020)
2. Atkins, A., Niranjan, M., Gerding, E.: Financial news predicts stock market volatility better than close price. J. Fin. Data Sci. **4**(2), 120–137 (2018)
3. Baron, M., Brogaard, J., Hagströmer, B., Kirilenko, A.: Risk and return in high-frequency trading. J. Fin. Quant. Anal. **54**(3), 993–1024 (2019)
4. Bartoletti, M., Pes, B., Serusi, S.: Data mining for detecting bitcoin Ponzi schemes. In: The 1st Crypto Valley Conference on Blockchain Technology (CVCBT 2018), pp. 75–84. IEEE (2018)
5. Bellia, M., Christensen, K., Kolokolov, A., Pelizzon, L., Renò, R.: High-frequency trading during flash crashes: walk of fame or hall of shame? SAFE Working Paper (2020)

6. Bendiksen, C., Gibbons, S.: The bitcoin mining network - trends, composition, average creation cost, electricity consumption & sources. CoinShares Research, Whitepaper, December 2019
7. Bollen, J., Mao, H., Zeng, X.: Twitter mood predicts the stock market. J. Comput. Sci. **2**(1), 1–8 (2011)
8. Brown, D.R.: Generic groups, collision resistance, and ECDSA. Des. Codes Crypt. **35**(1), 119–152 (2005)
9. Brumley, B.B., Tuveri, N.: Remote timing attacks are still practical. In: Atluri, V., Diaz, C. (eds.) ESORICS 2011. LNCS, vol. 6879, pp. 355–371. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23822-2_20
10. Budish, E., Cramton, P., Shim, J.: The high-frequency trading arms race: frequent batch auctions as a market design response. Q. J. Econ. **130**(4), 1547–1621 (2015)
11. Bujari, A., Furini, M., Laina, N.: On using cashtags to predict companies stock trends. In: Proceedings of the 14th IEEE Annual Consumer Communications & Networking Conference (CCNC 2017), pp. 25–28. IEEE (2017)
12. Chavoshi, N., Hamooni, H., Mueen, A.: DeBot: Twitter bot detection via warped correlation. In: The 16th International Conference on Data Mining (ICDM 2016), pp. 817–822. IEEE (2016)
13. Chesterman, S.: 'Move fast and break things': law, technology, and the problem of speed. NUS Law Working Paper (2020)
14. Cresci, S.: A decade of social bot detection. Commun. ACM **63**(10), 72–83 (2020)
15. Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., Tesconi, M.: Social fingerprinting: detection of spambot groups through DNA-inspired behavioral modeling. IEEE Trans. Dependable Secure Comput. **15**(4), 561–576 (2017)
16. Cresci, S., Lillo, F., Regoli, D., Tardelli, S., Tesconi, M.: $FAKE: evidence of spam and bot activity in stock microblogs on Twitter. In: The 12th International AAAI Conference on Web and Social Media (ICWSM 2018), pp. 580–583. AAAI (2018)
17. Cresci, S., Lillo, F., Regoli, D., Tardelli, S., Tesconi, M.: Cashtag piggybacking: uncovering spam and bot activity in stock microblogs on twitter. ACM Trans. Web (TWEB) **13**(2), 1–27 (2019)
18. Cresci, S., Petrocchi, M., Spognardi, A., Tognazzi, S.: From reaction to proaction: Unexplored ways to the detection of evolving spambots. In: Companion Proceedings of the Web Conference 2018 (WWW 2018), pp. 1469–1470 (2018)
19. Cresci, S., Petrocchi, M., Spognardi, A., Tognazzi, S.: Better safe than sorry: an adversarial approach to improve social bot detection. In: The 11th ACM Conference on Web Science (WebSci 2019), pp. 47–56 (2019)
20. Da San Martino, G., Cresci, S., Barrón-Cedeño, A., Yu, S., Di Pietro, R., Nakov, P.: A survey on computational propaganda detection. In: The 29th International Joint Conference on Artificial Intelligence (IJCAI 2020), pp. 4826–4832 (2020)
21. Das, S.R.: The future of fintech. Financ. Manage. **48**(4), 981–1007 (2019)
22. Davis, M., Kumiega, A., Van Vliet, B.: Ethics, finance, and automation: a preliminary survey of problems in high frequency trading. Sci. Eng. Ethics **19**(3), 851–874 (2013)
23. Di Pietro, R., Caprolu, M., Raponi, S.: Next generation information warfare: rationales, scenarios, threats, and open issues. In: Mori, P., Furnell, S., Camp, O. (eds.) ICISSP 2019. CCIS, vol. 1221, pp. 24–47. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-49443-8_2
24. Di Pietro, R., Raponi, S., Caprolu, M., Cresci, S.: New Dimensions of Information Warfare, pp. 1–4. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-60618-3_1

25. Ding, X., Zhang, Y., Liu, T., Duan, J.: Deep learning for event-driven stock prediction. In: The 24th International Joint Conference on Artificial Intelligence (IJCAI 2015) (2015)
26. Evans, L., Owda, M., Crockett, K., Vilas, A.F.: Credibility assessment of financial stock tweets. Expert Syst. Appl. **168**, 114351 (2020)
27. Ferrara, E.: Manipulation and abuse on social media. ACM SIGWEB Newsl. (Spring), 1–9 (2015)
28. Galeshchuk, S.: Technological bias at the exchange rate market. Intell. Syst. Account. Fin. Manage. **24**(2–3), 80–86 (2017)
29. Glenski, M., Saldanha, E., Volkova, S.: Characterizing speed and scale of cryptocurrency discussion spread on reddit. In: The 28th International Conference on World Wide Web (WWW 2019), pp. 560–570 (2019)
30. Grahl, J., Lysandrou, P.: The European commission's proposal for a financial transactions tax: a critical assessment. JCMS J. Common Market Stud. **52**(2), 234–249 (2014)
31. Hasbrouck, J., Saar, G.: Low-latency trading. J. Financial Mark. **16**(4), 646–679 (2013)
32. Hendershott, T., Riordan, R.: Algorithmic trading and the market for liquidity. J. Financial Quant. Anal. **48**(4), 1001–1024 (2013)
33. Jiang, M., Cui, P., Beutel, A., Faloutsos, C., Yang, S.: Inferring lockstep behavior from connectivity pattern in large graphs. Knowl. Inf. Syst. **48**(2), 399–428 (2016)
34. Kharratzadeh, M., Coates, M.: Weblog analysis for predicting correlations in stock price evolutions. In: The 6th International Conference on Web and Social Media (ICWSM 2012). AAAI (2012)
35. Kushner, D.: Sony vs. the hackers. IEEE Spectr. **48**(5), 16 (2011)
36. Linton, O., Mahmoodzadeh, S.: Implications of high-frequency trading for security markets. Ann. Rev. Econ. **10**, 237–259 (2018)
37. Lo, A.W.: The adaptive markets hypothesis. J. Portfolio Manage. **30**(5), 15–29 (2004)
38. Mazza, M., Cresci, S., Avvenuti, M., Quattrociocchi, W., Tesconi, M.: RTbust: exploiting temporal patterns for botnet detection on twitter. In: The 11th International Conference on Web Science (WebSci 2019), pp. 183–192. ACM (2019)
39. Mendoza, M., Tesconi, M., Cresci, S.: Bots in social and interaction networks: detection and impact estimation. ACM Trans. Inf. Syst. (TOIS) **39**(1), 1–32 (2020)
40. Menkveld, A.J.: High frequency trading and the new market makers. J. Financial Mark. **16**(4), 712–740 (2013)
41. Michaelis, K., Meyer, C., Schwenk, J.: Randomly failed! the state of randomness in current java implementations. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 129–144. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36095-4_9
42. Mirtaheri, M., Abu-El-Haija, S., Morstatter, F., Steeg, G.V., Galstyan, A.: Identifying and analyzing cryptocurrency manipulations in social media. arXiv preprint arXiv:1902.03110 (2019)
43. Monaco, E.: What FinTech can learn from high-frequency trading: economic consequences, open issues and future of corporate disclosure. In: Lynn, T., Mooney, J.G., Rosati, P., Cummins, M. (eds.) Disrupting Finance. PSDBET, pp. 51–70. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-02330-0_4
44. Nassirtoussi, A.K., Aghabozorgi, S., Wah, T.Y., Ngo, D.C.L.: Text mining for market prediction: a systematic review. Expert Syst. Appl. **41**(16), 7653–7670 (2014)
45. Neyret, A.: Stock market cybercrime. Technical report, Autorité des Marchés Financiers (AMF) (2020)

46. Nizzoli, L., Tardelli, S., Avvenuti, M., Cresci, S., Tesconi, M.: Coordinated behavior on social media in 2019 UK general election. arXiv preprint arXiv:2008.08370 (2020)
47. Nizzoli, L., Tardelli, S., Avvenuti, M., Cresci, S., Tesconi, M., Ferrara, E.: Charting the landscape of online cryptocurrency manipulation. IEEE Access **8**, 113230–113245 (2020)
48. Pacheco, D., Hui, P.M., Torres-Lugo, C., Truong, B.T., Flammini, A., Menczer, F.: Uncovering coordinated networks on social media. arXiv preprint arXiv:2001.05658 (2020)
49. Rajesh, N., Gandy, L.: CashTagNN: using sentiment of tweets with CashTags to predict stock market prices. In: The 11th International Conference on Intelligent Systems: Theories and Applications (SITA 2016), pp. 1–4. IEEE (2016)
50. Schumaker, R.P., Chen, H.: Textual analysis of stock market prediction using breaking financial news: the AZFin text system. ACM Trans. Inf. Syst. (TOIS) **27**(2), 1–19 (2009)
51. Sornette, D., von der Becke, S.: Crashes and high frequency trading: an evaluation of risks posed by high-speed algorithmic trading. The Future of Computer Trading in Financial Markets (2011)
52. Tardelli, S., Avvenuti, M., Tesconi, M., Cresci, S.: Characterizing social bots spreading financial disinformation. In: Meiselwitz, G. (ed.) HCII 2020. LNCS, vol. 12194, pp. 376–392. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-49570-1_26
53. Vasek, M., Moore, T.: Analyzing the Bitcoin Ponzi scheme ecosystem. In: Zohar, A., et al. (eds.) FC 2018. LNCS, vol. 10958, pp. 101–112. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-662-58820-8_8
54. Voukelatou, V., et al.: Measuring objective and subjective well-being: dimensions and data sources. Int. J. Data Sci. Anal. 1–31 (2020)
55. Wu, B., Liu, L., Yang, Y., Zheng, K., Wang, X.: Using improved conditional generative adversarial networks to detect social bots on twitter. IEEE Access **8**, 36664–36680 (2020)
56. Xu, J., Livshits, B.: The anatomy of a cryptocurrency pump-and-dump scheme. In: The 28th USENIX Security Symposium (SEC 2019), pp. 1609–1625 (2019)
57. Zaborovskaya, A., Zaborovskiy, V., Pletnev, K.: Possibilities of preventing manipulative transactions on the stock market in the conditions of new industrialization. In: The 2nd International Scientific Conference on New Industrialization: Global, National, Regional Dimension (SICNI 2018), pp. 154–160. Atlantis Press (2019)
58. Zhou, X., Zafarani, R.: A survey of fake news: fundamental theories, detection methods, and opportunities. ACM Comput. Surv. (CSUR) **53**(5), 1–40 (2020)