IMT School for Advanced Studies, Lucca
Lucca, Italy

Arbitrage in the Bitcoin ecosystem: an investigation of the Mt. Gox exchange platform

PhD Program in PhD in Institutions, Markets and Technologies - Curriculum in Economics, Management, and Data Science

XXXII Cycle

By

Pietro Saggese

2021

The dissertation of Pietro Saggese is approved.

PhD Program Coordinator: Massimo Riccaboni, IMT School for Advanced Studies Lucca

Advisor: Prof. Angelo Facchini, IMT School For Advanced Studies Lucca

Co-Advisor: Prof. Nicola Dimitri, University of Siena

The dissertation of Pietro Saggese has been reviewed by:

Roberto Renò, University of Verona

Antonio Scala, CNR Consiglio Nazionale delle Ricerche

IMT School for Advanced Studies Lucca
2021

He who obeys his heart is equipped to order.

Contents

List of Figures

List of Tables

Acknowledgements

I feel deeply indebted to many people who have greatly inspired and supported me during my Ph.D. studies and the writing of this thesis. Completing this essay, a product of several years' work, has been a long and challenging path that, alone, I would have never been able to bring to a conclusion. First, I would like to express my gratitude to my advisors, Angelo Facchini and Nicola Dimitri, for helping and conducting me during all these years. Angelo supported me with guidance, encouragement, academic stimulus: I always appreciated his sincerity and frankness, and I am thankful for his effort in patiently supporting me throughout all these years. I had the chance to meet Nicola in the quality of Microeconomics and Game Theory Professor. I immediately loved his method of teaching, and his unique ability to instill the most complex topics with simplicity. He is for me an unrivaled model and point of reference. I am grateful to Nicola also for introducing me to Rainer Böhme, who taught me so much, even in the short time we spent working together during my visiting in Innsbruck, through his rigorous scientific approach and the dedicating spirit for work, as well as his invaluable wisdom, both as an academic and as a mentor. I wish to dedicate a special thanks to Alessandro Belmonte: his help was essential to convey my intuitions into rigorous econometric analysis, and our discussions were always illuminating. Besides being an excellent researcher, I enjoyed his openness to laughs and jokes.

Equally importantly, I wish to thank all the people who supported me throughout all these years and cherished my days, without whom every world of this essay would be void to me. First of all, my family: Francesco, Clementina, Daniela, Maurizio; the eldest, Edda, and the newcomer, Egle, witnesses of two different millennia; Irene, Valerio, Rebeca, Cristina, deservedly new members of the team.

Thanks, Elisabetta, for making me look at the future with a smile. Thanks to my friends of a lifetime, "i barellieri" Albi e Bazza, as well as Letizia and Agnese (let the IMT tradition go on!). Leonardo, Chiara and the beautiful Fedino, Francesca, Marta, Francesco M., Elisa: I know I'll always be able to count on you. Cento, Chiara, Nicola: are you shuffling

the deck of cards? Waiting for the next belote! Coralie, my gratitude goes inevitably to you, too. My VL friends Francesco, Jean Louis, Jean Marc, Maurice, Simon: always arguing on politics, always united by our legendary futsal matches and post-match dinners!

A wonderful group of new friends enriched my days in Lucca, and I am glad to have experienced this chapter of my life with them: Nicolò and Falco, the only and unique co-members of the Cools' Club, Sara, Anna, Costanza, Tatiana, Abhishek, Jacopo, Francesco, Yara, Ilaria, Laura. Stella, how many steps did we walk down from the bishop room, crying and laughing about it all?

And how many cakes and cappuccini for breakfast, as well as countless 'aperitivi stellari' did we all share at osteria Stellario? It would be impossible not to mention the place that was our second home for us all. Thanks to my teammates of the IMTeam, we shared many defeats and few wins together, but always as a group of good friends: Stefano, Matteo B., Silvio, Ivan, Hakan, Domagoj, Giacomo, Nilay, Jerry, Emiliano, Dimitris, Federico, Hamed, Luigi, Matteo S., Rodolfo, Virginio, Sean, Victor, Samuel, Deison, Rodolfo. Sara B., Roberto, Marta, Margherita, Alessia, Tiziano, I hope we'll have the chance to organize soon a tournament of ping pong or calcio balilla! Stefano, Bianca, and all the teammates of the Pallanuoto Lucca: I played with you what I discovered to be one of the most beautiful sports on earth. Sebastiano, I learnt a lot from you. I feel close to you and deeply esteem you for your moral integrity and ethic: work, dedication, perseverance, fraternity, loyalty.

Sandra, Svetlana, Cecilia, Michael, Max, Martin, Patrik (special mention for you - for rescuing me, by bike and under the rain, when I needed help): I know I have good friends in Innsbruck. Andrea C., your help and IT knowledge has been essential to me. Last but not least, a special thanks goes to Matteo C., for being the best IMT roommate I could ever ask for: his intriguing and challenging Odyssey across the America, started with a Fullbright scolarhsip and ended on the NBA courts, soothed our moments of sadness and filled our hearts with hope.

Vita

October 31, 1990	Born, Turin, Italy
Oct. 2013	Bachelor degree in Physics Final mark: 103/110 Università degli studi di Torino, Turin, Italy
Dec. 2015	Master degree in Physics of Complex Systems Final mark: 110/110 cum laude Università degli studi di Torino, Turin, Italy
Nov. 2016	Enrollment, Doctor of Philosophy in Economics, Management and Data Science IMT School for Advanced Studies Lucca
Feb-May 2019	Visiting student University of Innsbruck

Publications

1. P. Saggese, A. Belmonte, N. Dimitri, A. Facchini, and R. Böhme. "Arbitrageurs: who are you? Who, who, who, who? Identifying triangular arbitrage in the Mt. Gox ecosystem," *(working paper)*.
2. V. Dudjak, T. Alskaif, S. Khadem, D. Neves, A. Pena-Bello, P. Saggese, M. Andoni, M. Bertolini, Y. Zhou, B. Lormeteau, M. A. Mustafa, Y. Wang, C. Francis, F. Zobiri, D. Parra Mendoza, B. Bowler, A. Papaemmanouil. "Impact of local energy markets on the power systems: comprehensive review", *(working paper)*.

Presentations

1. P. Saggese, and R. Böhme, "Identifying the Arbitrageurs on Mt. Gox: First Insights from the Leaked Dataset," at *Harvard School of Engineering and Applied Sciences - Workshop on the Economics of information security (WEIS) 2019*, Boston, June 2019.

Abstract

The purpose of the thesis is to identify and describe the arbitrage activity conducted in the Bitcoin ecosystem at its early stages. This work is the first attempt in the literature to investigate empirically the individual behavior of the arbitrageurs, and provides evidence that they are few and sophisticated. I exploit a dataset containing the history of trades executed within the exchange platform Mt. Gox between 2011 and 2013. I follow and improve upon the established methods to preprocess the data by proposing a new approach whose validity is documented extensively. Crucially, trades are labelled with user specific identifiers, allowing to reconstruct the individual sequences of actions and thus to identify arbitrageurs, and explicit transaction costs are accounted for. The core of the work is thus the implementation of two novel methodologies that aim at identifying the triangular arbitrage activity within the Mt. Gox platform and the two-point arbitrage across Mt. Gox and two counterpart exchanges, Bitstamp and BTC-e. In the former I focus on the mispricings of the bitcoin price denominated in different fiat currencies; in the latter, I compare differences in price - across Mt. Gox and the counterpart exchanges - denominated in the same fiat currency. I classify as arbitrageurs respectively 23 and 49 users, for a total of 72. A comparison of aggregate statistics between arbitrageurs and non arbitrageurs is given and discussed. This work represents the first empirical contribution on arbitrage at the micro level that goes beyond anecdotal evidence: the findings challenge the textbook definition of arbitrage and demonstrate that arbitrage is conducted by a limited number of sophisticated and specialized investors.

Chapter 1

Introduction

Arbitrage is a founding and unifying concept in financial economics. In the simplest form, it is the simultaneous purchase and sale of the same asset in two markets at different prices. It leads to price convergence, a necessary condition for efficiency and equilibrium among markets, and in principle it yields riskless profits to the investors who conduct it. Yet, studies on practical arbitrage show that instead it is risky and costly, and market anomalies can arise in the form of persistent mispricings: as a matter of fact, limits to arbitrage do exist. Moreover, anecdotal evidence suggests that arbitrageurs are few and specialized, whilst the textbook definition of arbitrage envisages the existence of many small traders, each exploiting an infinitesimal fraction of risk. Nonetheless, the economic literature lacks of empirical studies providing an answer to the question: "*Who are the arbitrageurs?*". I investigate the market of the cryptocurrency bitcoin and seek an answer to these related questions.

Bitcoin is a communication protocol that provides the users with a peer-to-peer network to transfer the ownership of a digital currency - that is, bitcoin - without the need for an intermediary or a central authority. This protocol relies on (and is the first implemented application of) the Blockchain, an innovative technology that employs extensively cryptography. In its essence, the Bitcoin Blockchain is a distributed, public and decentralized ledger composed of cryptographically linked blocks storing the whole history of the bitcoin transactions; each participant records and maintains an individual copy, and the consensus on the state of the system is ensured by a mechanism called Proof of Work (PoW) without resorting to any third party.

The procedures envisaged in the original design to obtain and possess bitcoins are non-trivial and require the users to be direct participants of the network, and to understand well the Bitcoin design principles. To circumvent this obstacle, centralized services known as *cryptocurrency exchange platforms* provide interfaces to conventional payment systems, acting in practice as intermediaries, and allowing their users to trade units of cryptocurrency against fiat money. Typical exchanges manage and match orders in a private limit order book and update their customers' account balances in cryptocurrency or fiat money when trades are executed. Thus, trades on such exchanges are kept in a private ledger and have no effect on the public Blockchain ledger, unless users withdraw cryptocurrency from the exchange. Mt. Gox was the first relevant exchange platform service active in the Bitcoin ecosystem. The competitors Bitstamp, BTC China and BTC-e entered the market in mid-2011 and gained a considerable role especially from late 2012. Mt. Gox dominated the market from 2011 to 2014, with a market share above 80% for most of that time, before being filed for bankruptcy in 2014.

The purpose of my work is to identify and describe the arbitrage activity conducted in the Bitcoin ecosystem: such exchanges play a prominent role in this sense, as they are the place where demand and supply meet and the Bitcoin price formation occurs. They are thus the natural target for studying arbitrage in the Bitcoin ecosystem. Intuitively, one would expect the identification procedure to rely on a comparison of the bitcoin flows across exchanges, publicly observable on the Bitcoin Blockchain, with the related differences in published prices (thereby generating evidence of arbitrage). However, bitcoin transactions are too slow and (at times) too costly for this strategy: arbitrageurs maintain a stock of both bitcoins and fiat money in accounts at each exchange in order to react quickly to price differences. Thus, the public ledger does not provide any advantage to identify arbitrage transactions, and evidence of arbitrage must be searched within the private ledgers of the exchanges themselves.

Most importantly, the internal log of Mt. Gox, who was leaked to the public in 2014, is characterized by the essential feature that trades are labelled at the user level, and explicit transaction costs are reported. This allows to identify the sequence of actions of each trader, and consequently to analyse empirically both their trading behavior and the profitability of their strategies. The leaked dataset contains the history of trades executed within the platform in between 1 April 2011 and 30 November 2013: I thus focus on the investigation of the Bitcoin ecosystem at its early stages. My approach is based on exploiting information on the trades performed in Mt. Gox, and in two other major exchange platforms active at the time of the analysis, Bitstamp and BTC-e.

I mine the leaked dataset with two aims in mind: to investigate the magnitude of the triangular arbitrage activity within the Mt. Gox platform, and then to identify the two-point arbitrage activity across Mt. Gox and the two other relevant markets at the time in the USD market, that is, BTC-e and Bitstamp. My goal is to observe the dynamics of arbitrage at the micro (individual) level. To the best of my knowledge, this work is the first attempt in the literature that exploits user-specific information to identify exactly the arbitrage actions, describing their features, and outlining the trade patterns of the investors who conducted it. Rather than investigating aggregate data, my approach grounds on the reconstruction of individual sequences of actions and on their comparison to time series of bitcoin price differences - denominated against different fiat currencies within the same market, and denominated against the same fiat currencies across different markets. Eventually, I classify the Mt. Gox users into two distinct categories: non-arbitrageurs and arbitrageurs. The latter is further subdivided in two (partial overlapping) categories, that is, the expert and the non expert arbitrageurs. In my work I provide exhaustive evidence that specific features and characteristics distinguish each group.

The remainder of the thesis is structured as follows. In this Chapter I briefly describe the main fields of study at the base of my investigation, in order to provide a leading thread and background information key to understanding the purpose of my work as part of a broader context. I cover a series of interdisciplinary topics that range from the recent findings in computer science related to the Distributed Ledger Technologies (focusing on Bitcoin, which I discuss both as a digital currency and as a communication protocol in Section 1.1), to more traditional concepts of traditional finance, such as the concept of arbitrage and the microstructure of the financial markets (Section 1.2). Finally, Section 1.3 focuses on the intersection of the two fields and describes the working principles of the cryptocurrency exchange platforms.

Chapter 2 describes the data cleaning procedures implemented to pre-process the Mt. Gox dataset. This stage involves the deduplication and correction of misreported data; I follow the methods established in the literature and purposefully deviate to improve upon them. I report my results together with comparisons to aggregate sources of information to ensure the validity of the proposed method. Then I provide an overview of the main economic dynamics that characterise the Mt. Gox ecosystem, and I devise and fit a model to estimate the fees paid by the users in each trade. This model is based on a comparison between the transaction costs reported in the leaked dataset and the official fee scheduled posted by the same exchange.

The methodologies to identify the two types of arbitrage (and the investors who conducted it) are different: information required to detect the triangular arbitrage is entirely contained in the private ledger of a single exchange. I thus exploit in Chapter 3 the Mt. Gox dataset to identify the triangular arbitrage activity conducted within the platform. I implement an algorithm to detect the arbitrage actions based on the fact that the availability of user identifiers per trade allows us to observe the historical record of each investor. Consequently, I describe the arbitrageurs' behavior. A considerable difference appears between users that conducted arbitrage in a single or in multiple currency markets, as well as between those who conducted few or many actions. Similarly, strategies that involve splitting orders to reduce market impact, or entail the execution of non-aggressive trades only, are good indicators of expertise. Using these elements as a proxy for trade ability, I find that trades performed by non expert users are on average non profitable when transaction costs are included, while skilled investors conduct arbitrage at a positive and statistically significant premium. Finally, I exploit within-user (across hours and markets) variation and document that expert users make profits on arbitrage by reacting quickly to plausible exogenous variations on the official exchange rates. Based on these indicators, I classify only 23 users as expert arbitrageurs; however, they are responsible for the vast majority of the actions and profits, supporting the thesis that arbitrageurs are few and expert users.

To identify two-point activity, instead, it is necessary to compare information from two different exchanges, and the data from the two counterpart exchanges Bitstamp and BTC-e are less richer, as the trades from their published datasets are anonymized. In Chapter 4 I identify the investors who performed two-point arbitrage across Mt. Gox and the two counterpart exchanges Bitstamp and BTC-e. Exploiting again the availability of user identifiers, I reconstruct the sequence of actions of each user and match these sequences to 'ideal' sequences of arbitrage trades, considering the price differences between exchanges and a user-specific estimate of transaction costs. The subset of investors whose actual series matches the ideal series best are potential two-point arbitrageurs. I identify 1,441 potential arbitrageurs with all two counterpart exchanges. Then, I cross-compare the actions executed by the potential arbitrageurs in the Mt. Gox platform to the logs of anonymized trades from the counterpart exchanges, and I construct a second metric that indicates if matches (equivalent and simultaneous trades) are found. I post-filter the first estimation and further reduce the set of detected users to 10 arbitrageurs with Bitstamp and 45 with BTC-e (49 in total).

Chapter 5 provides a summary of the main findings. I compare the results from Chapter 3 and 4, which are consistent in providing evidence that arbitrage was indeed executed in Mt. Gox, and that such activity was conducted by a restricted group of sophisticated investors. Noteworthy, I highlight that no user conducted both triangular and two-point arbitrage, suggesting that arbitrageurs are also specialized users.

1.1 Bitcoin

Before introducing exhaustively Bitcoin, it is worth noting that the term is used, somehow misleadingly, to refer both to a digital currency and to the payment system itself: indeed, Bitcoin is a communication protocol that provides a public community of untrusted pseudonymous users with an online peer-to-peer network infrastructure, along with a set of rules encoding how to mint and transfer the ownership of digital valuable assets, that is, bitcoins[1]. Decentralization is the most important design feature and the reason why Bitcoin is believed to be a potentially disruptive technology: transactions are executed without the need for an intermediary, and once the consensus on the state of the system is reached, transactions are stored without resorting to any third party on a public ledger distributed across participants, the so-called Blockchain.

Thus, both statements are true: Bitcoin is a communication system that uses extensively cryptography to reach consensus across untrusted participants, and bitcoin is the first digital currency that enables peer-to-peer transactions with physical cash-like features, such as no need for an intermediary and anonymity.

Incidentally, it is often not clear to what extent Bitcoin and Blockchain are two disentangled concepts. In these paragraphs I try to clarify some of these elements and briefly describe both Bitcoin and the Blockchain (more broadly, the concept of Distributed Ledger Technology)[2]. To do so, I first describe Bitcoin as a digital currency, and subsequently its characteristics as a communication protocol.

1.1.1 Digital currencies

Decentralization is the most important feature of Bitcoin. Thus, I first provide a broader definition of this concept, so fundamental in this frame-

[1] An established convention in the literature is to refer to the protocol using the capital letter (Bitcoin), and to the currency in lowercase letters (bitcoin).

[2] The main references used to cover this brief introduction to Bitcoin are Nakamoto (2008a), Bonneau et al. (2015), Böhme et al. (2015), Narayanan et al. (2016), Halaburda, Sarvary, et al. (2016), and Antonopoulos (2017).

work, in the light of a remark by Buterin (2017)[3]. In his view, the concept of decentralization is often misconceived, as in many cases it is implicitly intended as an abstract, not clearly defined concept, or as a synonym or notion complementary to the term 'distributed'. According to Buterin, decentralization can unfold in three main dimensions: *architectural decentralization*, *political decentralization*, and *logical decentralization*. Whilst this classification is envisaged mainly for computer systems, it can be thought at a more abstract and broader level[4].

By **architectural decentralization** it is intended the physical structure of the system[5]: how many computers compose a system, and to what extent such system will resist to a shutdown of a fraction of the computers? In other words, does the system have any single point of failure? A peer-to-peer network can be considered a fully architectural decentralized structure, in the sense that each participant is totally independent from each other and the network does not depend on the participation of any of the single users. Bitcoin relies on a peer-to-peer network that has no central point of failure, nor depends on third parties. **Political decentralization** arises when it is not possible to identify an entity or institution who has a form of control on the system; that is, when the system is not the liability of an entity. In Bitcoin nobody is responsible for the upkeep of the network; rather, it is a public good with embedded economic incentives that allow for its maintenance without relying on a central authority. Finally, the meaning of **logical decentralization** is more subtle. A system is logically centralized when the parts of the system behave coherently, that is, when the system appears as a 'monolithic object'. E.g., it may be useful to think of law, a single corpus, equal for every person: 'the law is the law'. Buterin suggests, as a heuristic to

[3]Vitalik Buterin is the founder of Ethereum, that is, a cryptocurrency that grounds on a DLT technology similar to the one implemented in Bitcoin. Highlighting the differences between Bitcoin and Ethereum goes beyond my scope. Should it be of interest for the reader, see e.g. Buterin et al. (2014) and Antonopoulos and G. Wood (2018).

[4]An alternative and interesting discussion on the interpretation of decentralization is given in Allen et al., 2020. Decentralization can take place in the form of *role separation*, *trust dispersal*, and *threshold trust*.

[5]Roughly speaking, this facet of the term decentralization is comparable to what is usually meant by 'distributed'.

classify a system as logically centralized or decentralized, to 'split it in half': does it still work properly? Bitcoin is logically *centralized*: people agree and reach a consensus on the state of the system, that is, who owns what amount of the digital valuable asset bitcoin. It is impossible to split such network while preserving its functioning.

In the light of this definition, I first discuss the novel aspects brought by bitcoin as a currency and in contrast with other forms of money; I will also profit of this brief description to introduce the cryptographic protocols relevant for the Bitcoin system. Comparing bitcoin to other forms of money may be helpful to understand better where the innovative side of Bitcoin resides. To do so, I refer to Figure 1, which represents an interpretation and re-elaboration of the taxonomy of money introduced in Bech and Garratt (2017)[6] by including the contribution of Buterin on the meaning of decentralization.

Coins and banknotes - physical cash - offer two distinctive features. First, everybody can exchange them with everybody else, i.e., in a peer-to-peer fashion, without resorting to an intermediary[7]: no bank intervenes when you pay with coins your coffee at 1 a.m. in the library, to write your Ph.D. thesis in time. Second, such transactions are mostly anonymous, i.e., there is not an intrinsic way to keep trace of physical cash transactions. This combination of properties was common also to commodity money, that is, objects like metal or food who were used in past as money for their intrinsic value; interestingly, money in this form shares with Bitcoin the feature of not being the liability of an institution. On the other side, commodity money does not share the property of being logically centralized, while cash does; there is no necessarily agreement on the value of a commodity, while everyone agrees on the value of one dollar. Most importantly, it has been overwhelmingly hard to replicate the combination of these two features (anonymity and absence of intermediaries) for electronic cash - indeed, before Bitcoin offered a solu-

[6]see also CPMI (2015) and Bjerg (2017).

[7]However, intermediaries do play a role in the minting of physical cash: the amount in circulation, e.g., is controlled and regulated by a central authority, normally a central bank. This also ensures that cash is a logically centralized system: everybody agrees, e.g., on the value of dollars.

Figure 1: A revisited taxonomy of money.

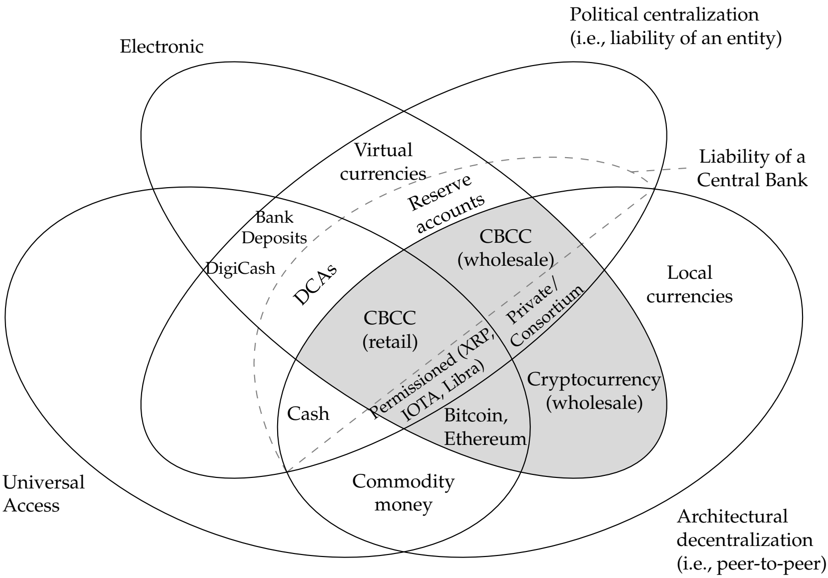

Notes: the classification is based on Bech and Garratt (2017) and Buterin (2017). I classify money according to the following dichotomies: 1) is universally accessible or local; 2) is physical or electronic; 3) is architecturally decentralized, that is, fully peer-to-peer or not; 4) is politically decentralized, that is, money which is 'not the liability of anyone', or governed by an entity or institution. The latter contains an important subset: money controlled by central banks. The gray area indicates a set of possibilities that were not technically feasible before the Bitcoin protocol was implemented. Bitcoin (as well as Ethereum) is public: everybody can participate to the peer-to-peer network without restrictions; a set of protocols (permissioned DLTs) slightly modify the Bitcoin original mechanism and propose an alternative scheme where some privileged users have formally a political role, while still providing a fully decentralized peer-to-peer network; private DLT protocols go even further and restrict the access to the network only to invited users, and are often used by private companies for internal accounting and data storage. See Ruiz (2020) for the classification of permissioned and private DLTs. Note that the distinction between these two areas is not always sharp, and depends on the individual projects and protocols. Finally, several projects of Central Bank CryptoCurrencies (CBCC) are currently in development.

tion, this had never been achieved. The gray area in Figure 1 identifies a set of typologies of currencies - that is, cryptocurrencies - based on the DLT technology and not technically feasible before the Bitcoin protocol was implemented.

Several attempts were done in past to reproduce the properties of physical cash in electronic format. A large fraction of the transfer of money as it is known today is indeed electronic: from central bank deposits to commercial bank money, the digitization is widely adopted for its clear advantages in simplifying transactions by reducing (time and speed related) costs. However, these are cases where the digital currency is an electronic representation of a physical counterpart regulated by an institutional entity. The idea of digital cash with no physical counterpart is relatively new and was introduced by the work of Chaum et al. (1983 and 1988), who implemented a cryptographic protocol called blind signature and subsequently commercialized his proposal under a company named DigiCash. This proposal is the first that applies digital signatures to a form of digital money.

Digital signatures are an essential concept in modern cryptography and are strictly related to two other cryptographic algorithms, that is, hash functions and public/private keys pairs. The three concepts are introduced in a seminal paper by Diffie and Hellman (1976); Rivest, Shamir, and Adleman (1978) provide a fundamental contribution to digital signature schemes by proposing the RSA algorithm.

Cryptographic hash functions are algorithms that take data inputs of arbitrary length and map them into strings of fixed length. Such functions satisfy some important properties: the output is deterministic, is computed efficiently (i.e., in short time), and every hash should be generated with the same probability; small changes in the output lead to large variations in the output; it is infeasible[8] to find two different inputs mapping to the same output (collision-free), and it is infeasible to invert the function, that is, to obtain the input knowing the output (hiding). Preneel (2010) provides an overview on existing hash functions.

Public-private keys are an inherently related cryptographic protocol

[8] It is plausible in theory, but not in practice.

based on one-way functions that take a private key as input (essentially, a code string that must be kept hidden to the public) and generate a public key, that is, a second string that is mathematically related to the private key; the one-way functions (e.g., elliptic curves) ensure that it is unfeasible to reconstruct the private key from the public one. The most important application of public-private keys are the digitally signed messages.

Like real signatures, digital signatures must satisfy the following properties: given the same message, only one private key can generate a certain signature, but everyone can easily verify its validity; and it must be tied to a given message, i.e., it must not be usable in other contexts. Thus, a digital signature is a function that takes as inputs a message and a private key and cryptographically provides a signature that does not reveal any information on the private key and the message. Any change in the two inputs will lead to a totally different signature. Moreover, everyone can easily verify it: a second function takes as input the message, the signature, and the public key, and exploits the mathematical relationship between public and private key to prove whether the private key associated to the provided public one was the one used to sign the message. The verification is straightforward, as the function is purposely not computationally intensive. A digital signature proves the identity of the sender of the message, the validity of the message, that a transaction is authorized by the sender, and that the content cannot be modified. Digital signatures can be used also by signing a message with the *public* key of a receiver, in order to send messages that only the receiver can open.

In Chaum's design, intended to facilitate transactions across users and merchants, the blind signature played an essential role as it was exploited both to provide anonymity for the users and to prevent double-spending, that is, the possibility to spend twice the same amount of money.

Double-spending is an issue of fundamental relevance for electronic money (as well as counterfeiting for physical money): indeed, electronic money is essentially a string of numbers, and in principle it could be copied as easily as it would be to copy and paste whatever digital string text, unless a measure to avoid such counterfeiting is implemented. A viable and simple solution would be to combine the digital string corresponding to money with a serial number and let a central authority verify that a transaction between two individuals is feasible, i.e., by assessing that the sender has not already spent such amount of money. (Note that this procedure is actually similar to how physical banknotes are traced and protected against counterfeiting.) The blind signature introduced by Chaum grounds on this intuition and improves upon it via cryptography by ensuring that the users can conduct transactions anonymously.

However, this mechanism suffers from some drawbacks: the system ultimately proposes a form of digital bank-issued cash relying on a central provider for transactions to be executed; rather than being a peer-to-peer system, it is a peer-to-merchant scheme; furthermore, information on the latter group is in part exposed to the central provider, who could infer their cash movements, and thus not fully anonymous. In this sense, DigiCash is not a good electronic substitute for physical cash. It filed for bankruptcy in 1998[9]. Chaum's proposal was anyway seminal, and many alternatives tried to improve some of its aspects, but none of them actually achieved a significant role in the market.

Other proposals which are relevant are the cost functions introduced in Dwork and Naor (1992), and Hashcash (Back, 2002). Scarcity is a prerequisite for a good to have value. DigiCash was inherently pegged to the value of the dollar and eventually was controlled by a bank; Back, Dwork and Naor introduced in two independent works a solution to fabricate scarcity for an otherwise un-metered (non limited) digital resource by imposing, as a precondition to access the resource, the execution of a costly computational task. Whilst they did not conceive the mechanism explicitly for monetary purposes, their work introduces implicitly a requirement which is mandatory for a digital currency to acquire value independently of other currencies[10].

Dwork and Naor devised an algorithm (pricing functions) as a solu-

[9] `https://www.forbes.com/forbes/1999/1101/6411390a.html`. This, and all the following links, were accessed on 18 June 2021.

[10] Unless an entity controls it: in this case, such digital resource would be metered, as the authority would control its provision; i.e., an authority controlling the digital resource and the design discussed here are two different ways to achieve scarcity.

tion to email spamming: as the cost of sending automatically the same mail to few or a million people is slightly increasing, malicious users have an economic incentive to perpetrate the spamming activity. However, if every sender had to pay a small enough amount of computational resources to perform the task, the additional cost would be negligible for normal users, but prohibitive for spammers, deterring them from executing it. An un-metered digital resource (e-mail spamming) is now limited by design. Adam essentially proposed independently the same mechanism, which he calls cost-function, acknowledging that this procedure is conceptually closely related to money minting: in his proposal he made explicit[11] the similitude with the minting process (not coincidentally, Satoshi Nakamoto[12] cites this work in his paper and not Dwork and Naor's proposal). However the intuition never evolved into a broader project designing a digital currency.

The idea of minting a digital currency by solving a computational puzzle was instead implemented in two other projects, namely Dai (1998) and Szabo (2005). However, ultimately, both proposals were not backed by implemented code and received limited attention. They are anyway important as they are closest to Bitcoin in that they try to solve the same problem: to provide a secure peer-to-peer framework for transactions among anonymous users, in order to make monetary transactions independent of governmental institutions or trusted third parties, by using a virtual currency not pegged to the value of any other physical currency (thus, scarce by design; noteworthy, as the same name suggests, Bitgold was chiefly a tentative to replicate a gold-like virtual resource). The two projects are inherently similar: first, in both cases, money is minted directly by solving the computational intensive task: the money supply is controlled by design by binding the minting process to the execution of a computational expensive action. Both rely on a peer-to-peer network

[11] "[...]In the context of cost-functions I use client to refer to the user who must compute a token using a cost-function MINT() which is used to create tokens to participate in a protocol with a server. I use the term mint for the cost-function because of the analogy between creating cost tokens and minting physical money.[...]", Back (2002) (p. 1).

[12] Satoshi Nakamoto is the alias used by the unknown inventor(s) of Bitcoin. See `https://nyti.ms/2TzEN1Y`.

of pseudonymous users[13], and both employ a secure timestamping protocol.

Secure timestamping was introduced by Haber and Stornetta (1990 and 1997) and refers to cryptographic algorithms intended to encrypt digital documents (more broadly, bit-string information) and simultaneously ensuring their chronological order of execution. Given a series of files incoming in chronological order, the idea is simply to hash the first file (bit-string) together with the date and a signature of the sender; the second step is to construct a hash pointer, that is, a data structure that contains a hash of the data and a string of code indicating (i.e., pointing) to where such information is stored, to retrieve it easily. The pointer will be attached to the second bit-string, together with the new date and signature, and hashed; and so on. Most importantly, not the date (order of creation) nor the content of any of the documents can be modified without altering all the subsequent files, as the manipulation of a single element of the chained documents (whether the content of the bit-string, the date, or the signature) would compromise the hashes of all the following documents, providing immediate evidence of tampering. Thus, secure timestamping offers a proof of the order of creation, guaranteeing simultaneously tamper-proof resistance. This simple yet powerful solution can be applied to more complex structures, e.g. Merkle trees.

Merkle (1980 and 1987) introduced a data structure that exploits cryptography to store large amounts of files in a single hashed string, by iteratively repeating hash functions on pairs of single (hashed) documents. E.g., consider 4 hashed files which I call $Hash_i$, i = 1,...4. In a simplified notation, a merkle tree is constructed as follows:

[13] Pseudonimity differs from complete anonymity as the identity is hidden through the use of pseudonyms; privacy is partially preserved, as the users are still exposed to some form of traceability.

$$Hash_{00} = hash(Hash_1 + Hash_2)$$
$$Hash_{01} = hash(Hash_3 + Hash_4)$$

$$Hash_{000} = hash(Hash_{00} + Hash_{01}))$$

The individual files are thus contained in a single string, $Hash_{000}$, that identifies the whole *block* of files in a compact form. The main benefit of merkle trees is that the number of computations necessary to assess the data integrity grows proportionally to the logarithm of the number of files stored (and not linearly). Applying the above described timestamping protocols to merkle trees instead of single documents is straightforward. The result is a *chain* of chronologically ordered *blocks* of files. (Such collection of timestamped blocks is, in simplified terms, the structure used in the Bitcoin protocol to store the bitcoin transactions: hence the name Blockchain.)

As stated above, the two proposals never took off in part because they were never implemented; however, they also had intrinsic limits: in b-money, users would broadcast transactions globally, and just a limited set of (trusted) participants would maintain and update the ledger of balances in accordance to such movements. Most of all, in case of disputes among parties, the author ultimately suggests to resort to an arbitrator. The design implies some form of cooperation among participants; it does not solve the double-spending issue between untrusted parties without resorting to an intermediary. Similarly, Szabo makes explicit that the minting and broadcasting should be performed "with minimal dependence on trusted third parties" (p. 1, Szabo, 2005): to record the transactions, users would rely on an external secure timestamping service. Still, third party is not utterly excluded.

Ultimately, both proposals rely somehow onto trusted servers. In addition, the authors realize that solving computational puzzles to enforce scarcity implies pegging the value of the minted money to the cost of computing power, but the protocols do not describe a detailed solution to manage the fact that computational costs might change due to, e.g., limited or inaccurate information provision, or technological evolution.

1.1.2 The Bitcoin protocol

Bitcoin succeeded where the previous projects had failed, that is, in finding a solution to the double-spending problem in a *purely* peer-to-peer network: the users can conduct value-based transactions using a digital currency, anonymously and without requiring the intervention of a trusted third party. A peculiarity of the Bitcoin protocol is that it does not introduce brand-new concepts or cyptographic functions; rather, it relies on several established methods[14]. The novelty regards how Satoshi Nakamoto (2008a) brilliantly merged them together to design a virtual currency embedding core decentralized technologies. The Bitcoin whitepaper is not an academic paper, and it does not provide all the specifications of the protocol; it is rather a description of the high-level idea, intended for an experienced and specialist audience, describing why the protocol should work. The reference point to understand Bitcoin design principles is the actual implementation of the protocol itself, which was developed by Satoshi Nakamoto between 2007 and 2009, and it is continuously updated by the Bitcoin community. The key components of the Bitcoin protocol are the following:

Users and the peer-to-peer network. An entity becomes a node of the Bitcoin network via a designated software[15] that generates pairs of public and private keys using standard cryptography. In principle, there is no way to recover from it information on the real identity of a person or institution[16]. The entities appear in the network as pseudonyms via

[14]It grounds on the aforementioned protocols: private-public keys to generate nodes to store and spend money on the network, digital signatures for cryptographic validation of transactions, hash pointers to timestamp blocks of transactions in Merkle root structures, and computational puzzles to secure them.

[15]Among the softwares that implement and follow the consensus rules, BitcoinCore is the reference open-source project `https://bitcoincore.org/en/about/`; for alternatives, see e.g. `https://bitcointalk.org/index.php?topic=4180898.0` and `https://bitcoin.eu/bitcoin-core-alternatives-dont-fork-blockchain/`

[16]However, many de-anonymization techniques indeed exist. Trivially, in a single direct

an *address*, obtained by hashing the public key, which is visible to every other node; it is used to publicly announce the existence on the network, and thus to receive bitcoins. The private key instead is hidden and is necessary to prove ownership of the bitcoins, and thus to redeem them.

The term wallet is used to identify the data structure that stores and manages the private keys under the control of the same node. The balance is the sum of bitcoins associated to such keys and thus redeemable by the user of a node. Interestingly, the protocol does not specify the concept of users or identities; rather, the Bitcoin network is composed of nodes that are essentially pairs of public-private keys. Similarly, also wallets and balances are superstructures not defined by the protocol and deducible by the history of transactions. (Note also that a user can manage more than one node.)

Ledger and Transactions. As stated above, Bitcoin is logically centralized in the sense that users agree on the state of the world, that is, on who can spend which bitcoins. Transactions are data structures used to broadcast on the network that two users are willing to accept a transition from a state s_0 to a new state s_1 where bitcoin ownership has changed: the output of a transaction is made of one or more digital strings, each identifying an amount of bitcoins and the new owner. The latter is not referenced directly; the string contains a code called *locking script* or *scriptPubKey* which is a cryptographic puzzle that only the new owner can decrypt using its private key.

Thus, at any point in time, the state of the world is represented by a series of transaction outputs that are unspent, in the sense that they are encrypted and redeemable only by the entity able to provide the correct private key unlocking them. Such strings of code are called unspent transaction outputs (UTXOs).

Each new transaction thus takes as inputs some of these UTXOs and generates *new* UTXOs that describe a new state of the system. Suppose

---

purchase, the merchant's identity can be associated to the used address; as transactions are publicly auditable, it is good practice to use a new key pair (and thus address) for each transaction, to avoid direct association to the other transactions executed. Transactions point by design to previous transactions: algorithms have been proposed to exploit this feature to link addresses and plausibly reconstruct the identities of users. De-anonymization can also involve the IP address used to broadcast information. See, e.g., Ron and Shamir (2013) and Meiklejohn et al. (2013) and the subsequent literature.

for simplicity that a new transaction is made of only one input and one output[17] . The procedure works as follows: the input contains a reference to the hash of the transaction that generated the UTXO that is being spent (i.e., the UTXO must be retrieved); and a *scriptSig*, that is, the part of the code that contains the solution to the puzzle in the referenced UTXO. Specifically, the scriptSig contains two elements: a digital signature executed by the sender (in this simplest case) on the referenced transaction and the hash of the public key of the new owner; and the public key of the sender. By signing the transaction, the sender ensures that it possesses the money and that the data in the signature cannot be modified.

The validation consists in processing the scriptSig and the scriptPubKey: the input calls by reference the UTXO generated in an older transaction; first, it verifies that the public key of the sender corresponds to the public key hash included in the referenced UTXO; then, it solves the computational puzzle by verifying the validity of the signature: indeed, a digital signature on a given message can be produced only by knowing the private key. On the contrary, it is computationally easy for everyone to take a digital signature, the message, and to verify whether a public key corresponds to the unique private key that generated such signature. If the one reported in the scriptSig confirms the signature, then the transaction is validated. This transaction produces a new UTXO containing an amount of bitcoins and the public key of the new owner included in a computational puzzle that allows only such entity to unlock it in further transactions.

The transactions proposing the new UTXOs (the new states of ownership) are broadcasted to all other participants. By construction, on average every ten minutes the overall state of the system is updated: a block made of pending transactions is collected and stored in a Merkle tree structure; then it is added to the chain containing all other previous blocks of transactions (hence the name Blockchain). Each block contains a hash pointer to the previous block, so that the tampering of one single

[17]This is the simplest case, which is the one of interest for this dissertation, and it is called pay-to-pub-key-hash (P2PKH). More complex transactions can involve more parties and thus more private keys.

transaction will modify the hashes of all the subsequent blocks. Besides being tamper resistant, the chain ensures also a timestamping ordering of the transactions. Note that each node stores an 'individual' blockchain that represents the personal view of the system; in addition, every user independently verifies and validates cryptographically all the transactions proposed in a block before accepting the block in their personal blockchain.

However, ultimately everybody agrees on how to update the system, that is, everybody ends up accepting or rejecting in a coordinated manner the incoming blocks: each blockchain, which is updated independently of each other, reflects the *same* view of the system. How is it achieved? And how is double spending prevented? Up to now, I described a protocol which is similar to Chaum's and the other examples reported above. However, their solution ultimately had to resort to an intermediary to secure no double-spending; in Bitcoin, there is no need of it, as trust in a third party is substituted by the trust that the computational power invested to store the blocks in the chain is large enough to prevent an attacker from tampering (by double spending) existing transactions.

Consensus and mining. The procedure to add a new block to the chain works as follows. Every user can decide (but is not forced) to participate to a 'competition' that consists in trying to solve a task which is resource intensive and thus costly: the first that finds a valid solution is entitled to update the ledger by proposing a new block of transactions to be added to the existing chain of blocks. This task is called Proof of Work (PoW), and is based on the idea of cost functions proposed by Dwork and Naor and Hashcash. Users that invest computational power in the 'competition' are called miners. In simple terms, every miner collects some transactions in a Merkle Tree structure and computes its hash[18]. Such hash is coupled with a parameter called nonce and hashed again. The goal of the Proof of Work (PoW) is to find a nonce such that the hash

[18]Note that each miner will likely construct different blocks: usually there are more pending transactions than available memory 'slots' in the Merkle tree; in addition, transactions are broadcasted on a global network: different miners can have access to different transactions at a given moment due to latencies.

starts at least with a certain number of zeros. This number is established by the target, a parameter setting the difficulty of the PoW: the higher the number of zeros required, the harder the PoW. The difficulty is variable and is based on the computational power invested in the mining, and is re-targeted when the last 2.016 blocks took on average more or less than 10 minutes to be mined. The solution can be found only by brute force, that is, by changing the nonce until the hash has enough zeros. Miners invest resources as they have economic incentives: by design, adding a block implies earning an amount of bitcoins.

The winner miner *proposes* and broadcasts to the network an updated version of the state of the system, which only differs in that it adds an additional block at the end of the chain. All the other users verify the new transactions, and they update their own blockchain to the proposed version only if all these transactions are correct.

Suppose instead that the block contains double spending. The network is momentarily divided: there are two competing versions of the chain, and there is no consensus, as the miner proposed a fraudulent block which is not accepted by the rest of the network[19]. The only possibility for the fraudulent node to let prevail its version of the blockchain is to have more computational power than the sum of all the other honest participants: once its block is added, a new competition begins. The malicious attacker (or attackers) would build on the new block, while the honest community would reject it and mine on the last accepted block. On average, the group of users that will impose its version is the one with more computational power, as the consensus is always achieved on the longest chain (that is, the chain that embeds the highest amount of computational power).

Thus, assuming that users are small and each one has only a fraction of computational power with respect to the total, the next competitions would be likely won by the larger group of honest users, who would

[19]Note that, should the task be *not* computationally intensive, there would be no deterrent for the users in the network to spam continuously malicious blocks with already spent coins, as the creation of such blocks would not be costly, exactly as the email spamming that Dwork and Naor tried to deter! The fact that the procedure is costly prevents users from spamming blocks embedding double spending.

append the new blocks on the penultimate (i.e., the last legit one for the honest community), ignoring the last block. The malicious user would be unable to spend the reward in bitcoins, mined on a 'dead branch' of the chain not accepted by the rest of the community. Should instead the malicious attacker(s) control more than 50% of the computational power in the network, then it would end up being able to impose its version of the chain to the honest participants, as in the long term it would be able to win more 'competitions', and thus eventually would propose the longest chain[20].

Novelties: In its essence, the innovation of the Bitcoin protocol lies in the mechanism envisaged to reach consensus across untrusted parties, without resorting to an intermediary, and to consequently store the record of bitcoin transactions in timestamped blocks secured by computational power; such process is executed independently (i.e., each node updates independently its chain), but eventually consensus on the state of the system is achieved.

Timestamping is not trusted, in the sense that it is not provided by third parties but by the miners, and ensures the ordering of transactions: actual timestamps are not essential, what matters is the determination of a relative ordering through the pointers from block to block. The timestamping is secured by the amount of computational power invested and implicitly embedded in the Blockchain: such computational effort prevents tampering of the stored transactions, as this would imply redoing all the proof-of-work embedded in the subsequent blocks, and double-spending, as the reward in bitcoins for mining a block is obtained only if all the included transactions are correct. It is precisely in this sense that

[20]In computer science, distributed computing refers to a network of independent components that communicate to coordinate their actions. Errors may arise, due to potentially faulty parts or an unreliable network. A fundamental problem for such distributed systems is to avoid the so-called Byzantine faults, that is, situations in which the system responds inconsistently to failure-detection, as the individual components do not succeed to reach consensus on the state of the system. The name of this class of errors comes from Lamport, Shostak, and Pease (1982), who described allegorically the faulty parts of a system as Byzantine Generals preparing a coordinated attack. It is important to understand to what extent a system is Byzantine fault tolerant, that is, to what extent it can provide a consistent outcome in the presence of faulty components. Bitcoin is a solution to the Byzantine Generals Problem (see Nakamoto (2008b) and Gramoli (2020)).

the introduction of the computational puzzle secures the blockchain: the network is secure up to the amount of (quantifiable) total computational power invested in the competition for the blocks, and up to how it is distributed across participants. The longest chain is the one that embeds the highest amount of computational power - and thus trust - from the network participants: as long as the majority is honest, the chain is secure. The trust in a third party is substituted by a protocol that deters misbehavior through economic incentives conditional to executing resource intensive tasks. Only indirectly the proof of work is used to mint the coins.

Interestingly, Bitcoin does not provide specifications on the entities, wallets, balances, or any other superstructure: in its essence, the data structure is nothing else than a database containing, in a public and easy to verify way, the historical record of the change of ownership of bitcoins, and a network of nodes having the rights to redeem them with private key cryptography. Each bitcoin is linked to its previous change of ownership, thus the full history of a bitcoin can be traced. The unique structure introduced is the chain of timestamped blocks that stores the transactions.

1.1.3 Blockchain and Distributed Ledger Technologies

Such structure, called Blockchain, is the backbone of the Bitcoin protocol and is an example (the first implementation) of a broader class of databases called Distributed Ledger Technologies (DLTs). The Blockchain is so innovative not as a data structure in se, but because it is combined with the idea to use proof of work to securely timestamp and store transactions without requiring a third party. Thus, DLT identifies whatever data structure - e.g., a Blockchain, the chain of blocks in Bitcoin; or Directed Acyclic Graphs in IOTA - exploited to record the history of the ownership of a digital valuable asset exchanged in an architecturally decentralized structure, but reaching a logically centralized state of the system, i.e., consensus.

The Blockchain has no central authority, that is, there is no need for a third party to execute the transactions. In addition, the network is composed of nodes with the same read and write access to the system (for instance, miners are not privileged users: every node deliberately decides whether to participate in the mining competition or not); in this sense, Bitcoin is politically decentralized[21]. Nonetheless, more recent cryptocurrencies introduced intermediate forms of political centralization by design, where e.g. only some nodes have write permission (permissioned DLTs), or participation is only on invitation (private or consortium DLTs): in that latter case, both write and read permissions are not granted to every user.

1.2 The microstructure of financial markets

This section deepens two related concepts: first, I introduce the micro level structure of the financial markets, and describe its main features. Second, I focus on the arbitrageurs, a specific category of informed market participants that ground their trading strategies on seeking price inconsistencies of financial instruments with respect to their fundamental value.

1.2.1 Market microstructure

Market microstructure is the field of research in finance that investigates how the financial markets operate, and how the micro level design and structure, as well as the internal frictions, affect the overall price and the market dynamics. Understanding their working principles is of fundamental relevance: the essential function of the financial system is to foster the efficient allocation of capital, facilitating the transfer of funds

[21] While this is true by design, in practice it is debatable: developers that decide how to improve and modify the protocols do have political power. Miners are strongly subject to concentration dynamics and consequently have oligopolistic behavior: to increase the probability of winning, individual miners use private services called mining pools that redirect their hashing power on a single node - and subsequently redistribute earnings based on the individual contribution. Currently, more than 50% of blocks are mined by less than 10 mining pools.

from participants with a surplus to those who are in shortage, and such transfers take place mainly in financial markets.

In essence, financial markets are (physical or virtual) dedicated facilities that enable and favor the purchase and sale of tradable securities - that is, claims on the borrower's future income or assets - in organized markets. Despite in practice their functioning is extremely complex[22], the general principles that determine the internal trading dynamics depend on the market design of the trading venues, which are common to the vast majority of such instruments.

Indeed, the financial instruments are generally issued in the primary markets, where securities are sold to other organizations directly by the issuer. This phase involves primarily investment banks that buy large quantities of securities and guarantee a price to the issuing entity. The instruments are consequently sold to the public in the secondary markets. The issuer acquires new funds only in this first negotiation, and the price paid on the primary market is based on the price that the investment banks expect will be set in the secondary markets through the buy and sell interactions among individual investors. Thus, as a matter of fact, the places where price formation occurs are the secondary markets. Moreover, such facilities allow greater flexibility as they make assets readily tradable, thus increasing their liquidity (the measure of the ability to easily trade an instrument and convert it into cash).

The participants of the secondary market are generally divided in two groups: the buy side, composed of institutional and individual investors who purchase trading services; and the sellers of such services, such as brokers and dealers. Brokers are middlemen that execute orders on behalf of clients[23], in a principal agency relationship; they earn profits by

[22]The instruments traded are countless and variegated, and many intermediaries, heavily regulated and with specific functions, play an essential role in channelling funds between parties. An overview of the financial system goes beyond the scope of this section; a broader description can be found in Mishkin (2018). In addition to this source, I based the following introduction on the works of Larry Harris (2003), De Jong and Rindi (2009), and Barry Johnson (2010).

[23]Currently, direct trading is restricted to members of the trading venues. To circumvent such restrictions, intermediaries can provide Direct Market Access (DMA) services to clients, by allowing them to use the broker's infrastructure and directly submit orders.

charging commissions on the executed trades. Dealers instead facilitate trades on behalf of themselves, that is, they operate on their account as *principals*, and possess capital. They earn profits by supplying liquidity to the market, i.e., they post information regarding the total volume they are willing to buy or sell and at what price, and they profit from the difference between the two.

Investors define their strategies by communicating a series of orders to their brokers, who consequently submit them to the trading venues. Orders are essentially instructions, or algorithms, that encode whether and how to buy or sell a given security in response to specific changes of the market conditions. Typically, the providers of exchange services offer to investors a wide range of possible executable orders, that allow to control several aspects of the trades (such as the desired price, the timing and/or lifetime, the behavior in case of partial execution, or more complex options to link trades and generate cascade effects among orders). The two most important types are market and **limit orders**. The latter is an instruction to execute a trade at a price not worse than the specified one. Limit orders provide liquidity to the market, in the sense that they stand ready to be filled when the desired market conditions are met. They never break the imposed price limit, thus they eliminate price risks, at the cost of increased timing and execution risks. **Market orders** are executed immediately, at the best (bid or ask) price on the market. They ensure speed and execution, at the cost of increased price risk: market orders are always completely filled, even if it means accepting matches with orders placed at unfavourable conditions (i.e., bids lower than the best bid or asks higher than the best ask). This also implies that they erode liquidity, contrary to limit orders. For these reasons market orders are classified as aggressive orders. Many orders are hybrid forms of the limit and market. Others may indicate further instructions that specify the behavior conditional to certain market conditions; e.g., stop orders are submitted only when the price reaches a threshold specified by the order. Hidden and iceberg orders instead are used by investors seeking to place large orders without affecting the price significantly, and allow them not to disclose publicly their full position. Orders can be linked together, so that the submission to the market is conditional to the execution of other orders.

Financial markets are the trading venues where such orders are executed, and are commonly distinguished into two broad categories, or a hybrid of the two: order-driven and quote-driven markets[24]. Trading venues organized in the first form are commonly known as exchanges. Brokers submit the traders' buy and sell orders, which are collected and ordered by the trading venue in an order book. An automated engine matches them according to a predefined set of rules that determine the execution precedence (generally, priority is based on best price and time order). The distance between the highest priced bid - the best bid - and the lowest priced offer - the best ask - determines the *bid ask spread*, a measure of the market liquidity. The orders that are submitted at prices far from the equilibrium passively wait the correct market conditions to be met before being executed. The total volume of such orders is a measure of the *depth* of a market: the impact of large trades on price is smaller in deep markets. In the case of order-driven markets, the price formation mechanism is directed by the actual posted orders. In quoted-driven markets, instead, dealers provide quotations at which they are willing to buy and sell a given amount of a security; the efficiency of the market is guaranteed in principle by the competition among dealers. Such trading mechanism is of the type Over-The-Counter, as brokers are instructed by their principals to accept or decline the dealers' published offers based on specific conditions. The latter type of market is more suitable for negotiation among parties.

Another important aspect of financial markets is the mechanism used in the trade execution process, as it has direct implications on the final price upon which parties reach consensus. For this reason, the process is also called price discovery, as it reveals the actual prices the agree-

[24]A broader overview would be much more complex; for instance, alternative trading venues such as Dark Pools and Electronic Communication networks (ECNs) are gaining more and more importance in financial markets. The latter are similar to exchanges, with the difference that they facilitate direct access to customers; Dark Pools are venues that provide confidentiality (at the cost of reduced transparency), and thus are often used to guarantee privacy between institutional investors executing large orders.

ment between parties is grounded upon. The three main mechanisms used are bilateral trading, continuous auctions, and call auctions. The first one is a one-to-one interaction between known parties, it is typical of negotiation-based venues, and facilitates customized deals. The latter two are both auction-based processes; in the continuous one, traders can submit their orders without any time constraint. The orders are added to the venue's internal queue list, and the matching mechanism fills the orders according to the predetermined set of rules that determine the execution precedence. In call auctions, instead, orders are collected for a predetermined period of time and executed together in a second moment. While the first is more flexible and guarantees immediacy, the latter reduces price volatility and increases the depth of the market.

Finally, transaction costs are another key aspect of financial markets. Besides explicit costs such as fees, commissions, and taxes, many implicit costs arise and are hard to evaluate for traders. For instance, high spreads imply high costs to traders seeking immediate liquidity; delay costs may arise due to changing market conditions between the decision and the execution timing; and the impact of an order on the market may cause significant price variations, especially in illiquid markets. A correct identification of such costs is essential for the implementation of a profitable strategy.

The research in market microstructure focuses on the actual trading processes in secondary markets, and on how the features described above affect prices and volumes traded. Indeed, traditional finance abstracted for long time from the actual trading mechanisms and depicted the markets as perfectly efficient and frictionless. Recent research shows that this is not the case, and mainly three fields have been investigated. First, authors studied the primary features of the market design and how they affect the trading dynamics; e.g., Amihud and Mendelson (1987) and (1991) showed that returns have different behavior in the opening and closing phases of the intraday market, due to the different execution methods, and that the trading mechanism impacts significantly prices. Glosten (1994) and Foucault (1999) introduced and discussed the economic principles of limit order books, and Christie and Schultz (1994) exploited a sample of stock quotations at Nasdaq in 1991 to show that anomalies in the distribution of spreads are plausibly explained by a strategic (collusive) behavior of the involved dealers. The second relevant topic is the process that leads from individual expectations to price formation through price discovery and trade execution; Grossman and Stiglitz (1980) provided a model describing how imperfect transmission of information affects price discovery and consequently the price formation mechanism. R. A. Wood, McInish, and Ord (1985) studied the price formation conditional to micro level characteristics of trades such as size, frequency, timing; similar investigations on the relationship between the market size and volatility are conducted by Jones, Kaul, and Lipson (1994). The third relevant field is related to the analysis of the transaction costs, their impact on returns, and the possible strategies to minimize them. Roll (1984) introduced an influential method to measure the effective bid-ask spread in efficient markets, and Stoll (2000) described a series of methods to identify and quantify frictions in trading.

In the recent years the market was revolutionized by the use of electronic and automated trading tools: trading techniques that exploit algorithms to implement complex predefined sets of instructions are more and more common, and are known as algorithmic trading (AT); high frequency trading (HFT) is a strictly related concept, and encompasses the trading strategies based on the execution of many small intraday trades that exploit low latency to profit. Latency is the elapsing time between decision and execution of an order. A large fraction of the current research in this field focuses on how these techniques changed the way financial markets operate: e.g. Hendershott, Jones, and Menkveld (2011) and Brogaard, Hendershott, and Riordan (2014) respectively show that algorithmic trading is informed and consequently narrows spreads, and that high frequency trading has a positive effect on market liquidity. A. Kirilenko et al. (2017) investigate the effect of electronic trading on an event, known as 'The Flash Crash', that took place in the E-mini S&P 500 stock index futures market on 6 May 2020. From a different perspective, Budish, Cramton, and Shim (2015) assert that low latencies, HFT and AT made the current market design outdated and that frequent batch auctions would be more suitable. The dynamics of HFT are studied also in cryptocurrency markets (Daian et al., 2019).

1.2.2 Arbitrageurs, and limits to arbitrage

Arbitrageurs are a specific typology of investors that attempt to profit by seeking relative mispricings between assets. Expertise and knowledge are essential requirements for them to profit. Indeed, traders are commonly divided between *informed* and *uninformed*: the former ground their decisions on correct information regarding the *true value*, or fundamental value, of the instrument, while the latter "trade on noise as if it were information" (Black, 1986, p. 529) and are also called noise traders. Informed traders exploit both private and public information, and have an essential role in the markets as their activity conveys information into prices. Their actions reflect their own evaluation of the fundamental value. Instead, by basing their actions on noise or incorrect information, uninformed traders provide liquidity and thus create a market at prices far from fundamental values. Thus, despite they are a source of inefficiency in markets, they are also vital for financial markets in the sense that they make trading possible (Larry Harris, 2003; De Jong and Rindi, 2009; Barry Johnson, 2010).

Besides exploiting information itself, informed traders compete on the collection and incorporation of new information to predict the price changes and to identify predictable patterns or systematic errors. Arbitrageurs are a specific category of informed traders who focus instead on relative price differences between assets that share equivalent (or similar) factors determining their fundamental value. Their activity is essential, as it ensures the *law of one price*; that is, equivalent assets must be priced identically. Indeed, they buy undervalued assets and sell those overvalued, thus pushing prices towards equilibrium and enforcing the price convergence by concurrently profiting from this activity. The textbook description of pure arbitrage, in its simplest case, is that of the simultaneous purchase and sale of the same (or equivalent) asset at different prices: it yields risk-free profits, as the investor exploits advantageously the mispricing and the position is immediately hedged, without net investment of capital (Bodie, Kane, and Marcus, 2018).

More specifically, the theoretical approach distinguishes between two different types of arbitrage (Ingersoll, 1987). Given a set of states $\mathcal{S} = \{1, ..., S\}$ and N securities, the payoff matrix A_{ij} defines the payoff in the state i for each security j. The portfolio $\mathbf{x} \in \mathbb{R}^N$ is a combination of such securities, whose vector of prices is defined by $\mathbf{q} \in \mathbb{R}^N$. A state-price vector is a vector $\boldsymbol{\psi}$ such that $A^T \cdot \boldsymbol{\psi} = \mathbf{q}$. The arbitrage of type I takes place when the final payout is non-negative and no initial investment is required: that is, $\mathbf{q} \cdot \mathbf{x} \leq 0$ and $A\mathbf{x} > 0$. Type II arbitrage can be exploited if the cost of a portfolio is negative (i.e., one receives money *today* to hold the portfolio) and the payoff is non-negative: $\mathbf{q} \cdot \mathbf{x} < 0$ and $A\mathbf{x} \geq 0$. The first implies that one allocation stochastically dominates the others, while in the latter case the market includes mispriced redundant assets (linear independence implies that $A\mathbf{x} = 0$ only if $\mathbf{x} = 0$).

The fundamental theorem of asset pricing (or FTAP, whose first version is described in the seminal works of S. Ross, 1976, and 1978) states that, in a market with S states and N securities as the one described above, no arbitrage opportunities arise *if and only if* the vector $\boldsymbol{\psi}$ exists and is consistent with the security prices $\mathbf{q}$. That is, for a portfolio with payoff $A\mathbf{x}$ which is uniquely priced by $\mathbf{q} \cdot \mathbf{x}$ (linear pricing rule), there is a strictly positive state-price vector $\boldsymbol{\psi}$ such that the relationship $A^T \cdot \boldsymbol{\psi} = \mathbf{q}$ holds (positive pricing rule). If it is not possible to find such a strictly positive vector $\boldsymbol{\psi}$, then there are arbitrage opportunities in the market. The proof is based on the Separating Hyperplane Theorem (see Cerny, 2009; Björk, 2009; Duffie, 2010 who provide additional references on the topic). The absence of arbitrage is also implied in a theoretical framework of rational individuals who optimize their choices and prefer more to less: Dybvig and S. Ross (1989) show that such assumption is equivalent to the no-arbitrage condition (and thus to the existence of positive linear prices; see also Varian, 1987).

The contribution of Ross is seminal for the whole body of literature on the theory of asset pricing. The subsequent research on arbitrage, primarily focused on the mathematical generalization of these concepts, investigates the relationship between the no arbitrage condition and the martingale theory. The seminal paper of Harrison and Kreps (1979) is the first that explicitly reformulates the principles of the fundamental theorem of asset pricing in terms of martingale measures. The authors provide a general theory of option pricing based on the stochastic modeling, and show that no arbitrage opportunities exist if and only if it is possible to define a probability measure under which the price is a martingale. Most importantly, the prices are also consistent with those of the economic equilibrium. Harrison and Pliska (1981) turn this principle into a theorem for finite probability spaces, and Kreps (1981) extends it to a more general case by introducing a stronger assumption (known as 'no free lunch'). Delbaen and Schachermayer (1994) and (1998) introduce the concept of 'no free lunch with vanishing risk' (NFLVR), which is a slightly stronger assumption of the no arbitrage condition, and extend the formalization to semi-martingales: they show that the fulfillment of the NFLVR condition is equivalent to the existence of a semi-martingale measure equivalent to the original one, and this is *essentially* equivalent to the absence of arbitrage opportunities (see also Delbaen and Schachermayer, 2006).

In summary, the theory on arbitrage is well established, and probably one of the most studied concepts in the financial literature. Interestingly, however, practical arbitrage appears to be far from the one described by the theoretical approach. Theoretical arbitrage yields risk-free profits, and is conducted by fully rational arbitrageurs with homogeneous expectations who do not need any endowment of capital to perform it. However, this seems not to be the case in real markets. Arbitrage can be essentially risk-free only when it involves mean-reverting assets (securities whose mean price is known in the long run). In most of the cases, the instruments are non-stationary and thus intrinsically risky (i.e., prices do not converge to a mean value in the long run). Furthermore, also in the case of theoretically riskless arbitrage, arbitrageurs can end up trading at worse than expected prices (*implementation risk*); aggressive orders entail *price execution risk*, as they demand liquidity and thus have a higher impact on price. Conversely, limit orders are subject to *execution risk*, that is, the trade (or worse, only part of it) might never be executed. Arbitrageurs face also the risk of a short-term trend reversion of one of the two assets (*basis risk*), and may misinterpret a change in fundamental values as the effect of noise trading (*model risk*). Prices can take longer than expected to converge, and arbitrageurs can be forced to modify their positions in suboptimal situations, due to capital requirements (*carrying costs risk*).

Risk is not the only concern for arbitrageurs: markets are not frictionless, and investors need to account for additional costs such as fees, capital requirements, entry, strategy implementation, non-instantaneous diffusion of information (R. Merton, 1987). Capital requirements for arbitrageurs in an agency relationship may represent an issue: Shleifer and Vishny (1997) propose a model where, in the presence of extreme price deviations, principals evaluate erroneously the arbitrageurs (the agents) by judging them on their current performance, and refuse to supply additional capital precisely when it would be most needed, de facto limiting the arbitrage itself. De Long et al. (1990) show that irrational noise traders can operate on a bias, and consequently introducing an additional risk that limits arbitrage activity. Thus, as a matter of fact, arbitrage in real markets is far from the one described in theoretical works: Lawrence Harris and Gurel (1986), Froot and Dabora (1999), and Lamont and R. H. Thaler (2003b) provide empirical evidence of price anomalies and persistent mispricings; Barberis and R. Thaler (2003) and Gromb and Vayanos (2010) provide a survey on the recent developments of the analysis of the limits to arbitrage.

Somehow surprisingly, despite the authors of the latter work highlight the importance of understanding who actually are arbitrageurs in practice, the empirical evidence describing how user-specific characteristics of arbitrage lacks completely in the literature. Practicioners acknowledge that many small investors, each bearing an infinitesimal portion of risk are not the ones that would conduct arbitrage in practice. Rather, arbitrageurs are few and competing informed users that exploit sophisticated trading strategies that need capital. This statement is however mainly anecdotal and not supported by any academic study.

Finally, the major typologies of arbitrage are:

Shipping (or two-point) arbitrage: two identical securities are traded in different markets. Moving the security between the markets can bound the arbitrage opportunities (due to the shipping costs), but often it is not necessary in practice. **Conversion** arbitrage is performed between securities with similar - but of a different type - risk, to hedge it. **Triangular** arbitrage exploits price differences for the same asset quoted in different currencies: arbitrage opportunities arise if the implied exchange rate differs from the original one. E.g., in the Forex market it is possible to compare the official exchange rate of two currencies A and B with the rate implied against a third currency C. **Spreads**: the arbitrage is conducted on securities essentially identical, except for one element (maturity, credit quality, ...). **Merger (or risk)** arbitrage takes place when the stock prices of two firms involved in a merger react inconsistently with respect to the terms and cost agreed for the acquisition. **Pairs trading** consists in the identification of correlated assets, depending on similar common valuation factors: profits arise when only one of the two responds to price changes of the common valuation factor (viceversa, losses incur if the changes are due to instrument-specific factors). **Statistical arbitrage** exploits advanced mathematical and statistical methods. It can be thought of as an evolution of the former, applied to a bundle of instruments; in essence, the arbitrageur tries to identify their common factors, and determine price inconsistencies within the bundle.

1.3 Cryptocurrency exchange platforms

This section focuses on the exchanges dedicated to the trading of cryptocurrencies against fiat currency that blossomed in the last decade. The trading mechanisms on which they rely upon are in most cases adpoted from traditional finance, with some key differences to adapt them to the cryptocurrency characterstics. The majority of cryptocurrency exchanges are organized as order-driven markets with two sided continuous auctions; contrary to traditional market design, investors trade directly among them without resorting to brokerage services, and the role of intermediaries is much less - if not at all - regulated. Nonetheless, as a matter of fact such exchanges are centralized entities acting as an intermediary between the Bitcoin network and the traditional finance system.

Indeed, as described in Section 1.1, the Bitcoin design originally envisages two primary ways to obtain bitcoins: by winning the competition to add new blocks to the chain (the mining process), and by selling goods for bitcoins (the transactions). In both cases, users must be active nodes of the network and possess sufficient technological skills to master and run a dedicated software. This aspect raises some issues, as the protocol itself exposes users to high risks, if mishandled (see Conti et al. (2018) for a broad review on Bitcoin related privacy and security issues): first, transactions are irreversible, and no authority controlling the network can invert them in the presence of operational errors (e.g., sending the wrong amount, or selecting the wrong receiver[25]). Second, the loss of private keys is a common issue, and entails the impossibility to redeem bitcoins controlled by such key; again, there is no way to resort to an authority to recover them. A study executed by Chainalysis estimates that in November 2017 between 2.8 million and 3.8 million bitcoins (on a total of 21 million) were likely lost, that is, nobody has the control of the private keys to redeem them[26]. Even in the absence of operational errors, users can also be subject to attacks from malicious entities: some, like DDoS attacks (Benjamin Johnson et al., 2014) or the 51% attack described in Section 1.1, are distinctive of the Bitcoin infrastructure; others are common crimes facilitated by bitcoins, such as scams and frauds, hacks and bitcoin-denominated ransoms, that flourished around the Bitcoin ecosystem (Vasek and Moore (2015), and https://bit.ly/37Y6WZc). In summary, handling bitcoins correctly is non-trivial and several barriers to entry exist: especially at the early stages, the adoption of Bitcoin was confined to a limited number of users with a good knowledge and understanding of Bitcoin design principles and mechanisms. Despite this, Bitcoin started attracting the attention of less expert investors as an in-

[25]Bitcoins sent to a wrong address: https://bit.ly/3jRqCQJ , or https://bit.ly/3oMSBVm; wrong amount indicated in the transaction: https://bit.ly/3jJ2GyY.

[26]See https://bit.ly/32196Dl and https://cnb.cx/34NvcLn.

vestment asset, rather than an alternative payment system (Glaser et al., 2014; Yermack, 2015).

Exchange platforms came into play exactly for this reason: to reduce technological and switching costs for less expert users willing to possess bitcoins, and to facilitate trading against conventional payment systems for miners willing to sell their bitcoins in exchange of traditional fiat currencies. Indeed, the demand for bitcoins and the number of users relying on such platforms increased steadily from 2011 on (see e.g. Figure 18 in Section 3.2). Besides this, exchanges grew in importance also as they exploit economies of scale, by reducing the fixed costs a user would pay[27].

Currency exchanges are trading platforms that provide users with an online interface to trade bitcoins (and other cryptocurrencies) against fiat currencies. They share many working principles with traditional equity markets, where traders submit buy and sell orders and the exchange clears trades, providing the service for a fee. The first exchanges operated on single currency markets against bitcoins; e.g., Mt. Gox, the first relevant platform, at first allowed users to trade bitcoins against USD, and only subsequently additional fiat currency markets were introduced. In the last years institutional investors have entered into the market, the use of cryptocurrencies has become more widespread, and competition across exchanges has increased: while the first platforms where more rudimentary (though they already provided API structures allowing to conduct automated and complex order strategies), such services are now much more sophisticated and list often several crypto- and fiat- currencies, allowing to trade cryptocurrencies against all of them. Often exchanges focus on specific geographical locations, that is, list one or few specific fiat currencies and allow to trade it in a limited geographical area against several cryptocurrencies. Noteworthy, a subgroup of the current exchanges allows only trading among cryptocurrencies. Bearing this in mind, for the purpose of this dissertation I now focus on the dynamics regarding only the bitcoin markets against fiat currencies.

A key difference with traditional trading platforms is how funds are

[27]Costs related to security measures, to investments in knowledge and technology, to run the Bitcoin software on a personal device, ...

deposited into the exchanges: to start selling bitcoins, users must send them 'blindly' to a public key indicated by the exchange[28]; users with no bitcoins can enter the market by blindly sending an amount of money to the exchange that will buy for them bitcoins from the network and open an account at their name (noteworthy: not as a node of the Bitcoin network, but as a user of the online platform). To non expert investors, the platform simply appears as a web interface with a password and an identifier that gives access to a personal area reporting the user's balances, and an online market where to buy and sell fiat- and crypto- currencies. There is no need for them to know that exchanges are actually centralised superstructures basically acting as intermediaries between customers.

Under the scenes, indeed, exchanges are managed by a company that directly controls the customers' funds, through a series of bank deposits holding customers' fiat money, and at least one node of the Bitcoin network storing the bitcoins[29]. They control the keys, thus they ultimately possess the funds, and they appear in the Bitcoin network as very wealthy entities managing large amounts of bitcoins. On the (poorly developed) regulatory side, it is often not clear whether exchanges guarantee to just hold coins on behalf of the customers, as if they were gold merchants keeping labeled bars of gold - in this case, every trade across users would be executed and then transcribed on the public ledger - or if customers have a claim on an amount of money that the exchange commits to pay under request. In practice they always implement the latter, behaving as a matter of fact similarly to traditional banks (Anderson et al., 2018)[30]. This has two important consequences: first, they do not need to possess the exact amount of bitcoins owed at any moment in time; and second, they keep the balances updated off-line, on a private ledger:

[28]The private key is managed by the platform itself; nobody guarantees that the owners will not steal the bitcoins or use them for other purposes. Indeed, reputation systems are essential for the exchanges to be trusted and for the market to hold.

[29]Usually, the same exchange uses *cold wallets* with the highest security standards to store offline the vast majority of funds, and keeps the liquidity necessary for daily withdrawal/deposit operations in the so-called *hot wallets*.

[30]Decentralised exchanges (DEXs) represent an alternative business model, in which the exchange does not take the cryptocurrency in custody, and trades are executed via smart contracts on a blockchain.

trades executed in cryptocurrency markets are not stored on the Bitcoin Blockchain.

Figure 2 reports the evolution in time of the main exchange platforms[31]. The upper panel shows the evolution of the daily volume percentage traded in the main exchanges, while the lower one reports the percentage per currency. Time series are reported as a mean rolling average over three days. Three different temporal windows with distinctive patterns can be clearly identified: the first one, from 2010 to the end of 2013, corresponds to an early phase dominated by a single platform, that is, Mt. Gox; the vast majority of the trades executed in this platform are in USD. The exchange is notorious for a long story of security breaches (the first one caused the variations visible in both panels at the end of June 2011, when a security breach pushed the bitcoin price to almost 0$ in few hours, and forced the platform to stop operations for some days); it ceased activity on 7 February 2014, when it halted all withdrawals and was subsequently filed for bankruptcy; most of the customers sued the exchange in order to recover their funds, and the case is still open. This platform is especially important for my work, as I exploit the private ledger containing the log of internal trades (that was leaked and published in 2014) to conduct the analyses.

The second time window is dominated by trades denominated in CNY, and most of the orders were cleared in three major platforms: OKCoin, Huobi, and BTC China. This phase ended in 2017: alarmed by the growing use of bitcoin to circumvent capital controls, the People's Bank of China's (PBoC) intervened directly in January 2017, and pressured the exchanges into changing their business models by introducing fees and halting margin trading; the volumes in these platforms felt abruptly in few days, and eventually in September the Chinese government ordered to completely shut down the country's bitcoin exchanges[32].

The last (and current) phase is dominated by few but competing platforms with a solid reputation system, and trading is conducted mainly

[31] Data from `https://data.bitcoinity.org/`. The Mt. Gox data are based on personal estimations obtained by the exchange's private ledger.

[32] See `https://on.ft.com/3820Uqu`, `https://bit.ly/2HZwqe4`, and `https://on.wsj.com/3kWkLuJ`

Figure 2: Evolution in time of the cryptocurrency exchanges trading dynamics.

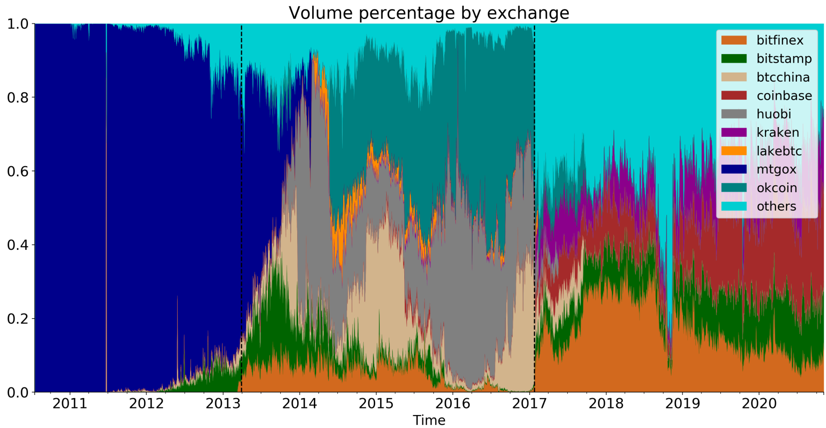

(a)

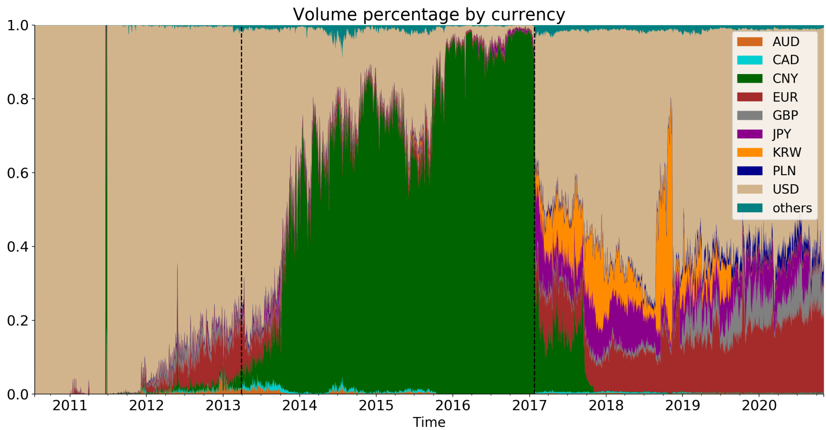

(b)

Notes: evolution of the daily volume percentage traded in the main exchanges (Panel a), and percentage per currency (Panel b). Three different phases with distinctive patterns appear. The first coincides with the preponderant presence of Mt. Gox, from 2010 to 2013; the second, from 2014 to 2017, corresponds to a time window mostly dominated by Chinese exchanges; the last and current phase is more heterogeneous, and few but competing exchanges share the market and operate against multiple fiat currencies.

Figure 3: Fee scheme for the main exchanges active in the USD market between 2011 and 2013.

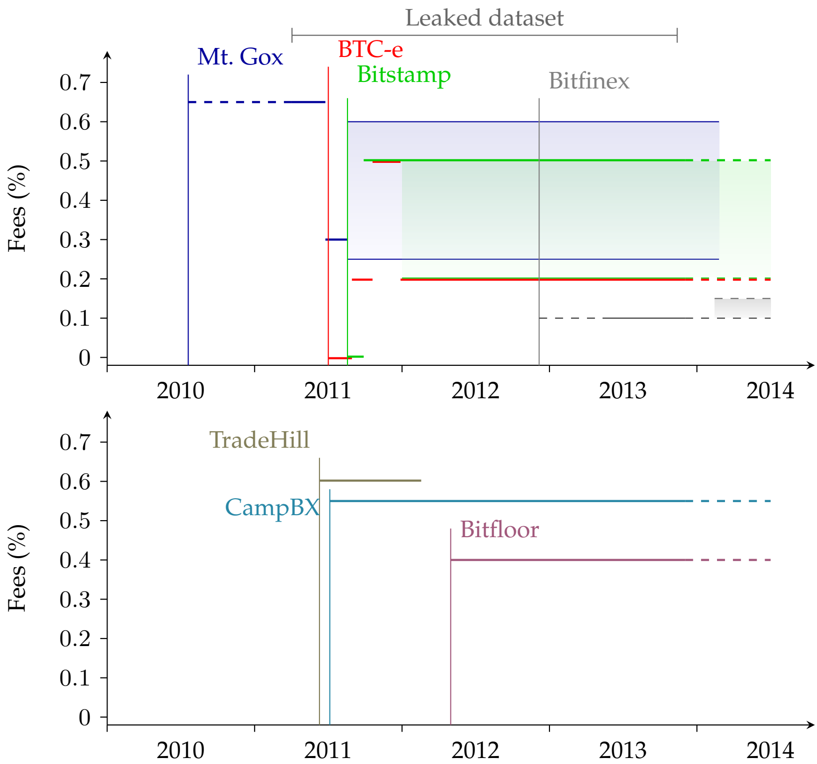

Notes: The plots show a comparison in time of the fee schedule across the main exchanges active at the early stages of the Bitcoin ecosystem. Dashed lines refer to time windows where I inferred the fee scheme, or that go beyond the time window covered by the Mt. Gox leaked dataset. Mt. Gox transaction costs went from 0.65% to 0.3% in June 2011, and from mid August the exchange introduced a scheme with fees varying depending on the volume traded in the last month. Bitstamp adopted a similar schedule from January 2012, after imposing a fixed fee for short periods of time (0% and 0.5% between mid June and the end of 2011). The transaction costs in BTC-e followed a similar pattern with respect to Bitstamp (0%, then 0.2% and finally 0.5% between June and December 2011). The fees finally stabilized at a fixed 0.2% rate. Bitfinex, who entered later in the market, imposed instead a fixed 0.1% fee and later introduced a distinction between market makers and takers (up to 0.15%). The lower panel reports other exchanges, that is, TradeHill, CampBX, Bitfloor: each of them offered stable fees over time (respectively 0.6%, 0.55%, and 0.4%).

in USD, but a relevant market share is executed in other currencies like EUR, GBP, JPY. CoinmarketCap and Bitcoin Wiki[33] provide a detailed list of the most prominent exchange platforms active to date.

Figure 3 focuses on the first time epoch identified, which is the most relevant for the following chapters, and reports in two separated plots the evoluton in time (x-axis) of the percentage fees (y-axis) that users would pay in the main exchanges active in the USD market before 2014, that is, Mt. Gox, BTC-e, Bitstamp, Bitfinex, TradeHill, CampBX, Bitfloor. As one can see, the patterns vary greatly across them: some venues modified their fee scheme often in time, and offered discounts to users with high trading volume, such as Mt. Gox; others, such as CampBX, or Bitfloor, offered instead stable fees over time. Overall, the comparison shows that fees tended to decrease in time.

Whilst in the recent years a large number of alternative exchanges were launched, it is widely acknowledged that most of them are scams or report fake volumes to lure customers in; a study by the BitWise Asset Management Company (2019) proposes a few methods to spot fabricated data and show that only a small number of the listed exchanges are indeed working transaparently and can be trusted[34]. When the market is narrowed down to these few companies, it appears more ordered, efficient, much smaller in terms of volumes traded, and competitive across exchanges for the services offered. This aspect highlights one of the most prominent risks associated to trading on such platforms: exchanges can host (or be) fraudulent entities that manipulate markets in order to send misleading signals to customers. Furthermore, as the textbook example of Mt. Gox shows, exchanges are also frequently subject to shutdowns and most of them do not refund customers when ceasing operations[35]; see Moore and Christin (2013) for this form of counterparty risk. For completeness, another important risk associated with the exchange plat-

[33] Respectively `https://coinmarketcap.com/rankings/exchanges/` and `https://en.bitcoin.it/wiki/Comparison_of_exchanges`.

[34] E.g., some of the major ones are Binance, Bitfinex, Kraken, Bitstamp, Coinbase, Bitflyer, Gemini, itBit, Bittrex, Poloniex.

[35] Or the owner of the exchange might die without a secure back up of the private key: that really happened to QuadrigaCX, in December 2018 (`https://bit.ly/2JCzo8T`).

forms is that of price manipulation: the lack of regulation facilitated fraudulent users to conduct trading techniques otherwise illegal in traditional financial markets (Gandal, Hamrick, et al., 2018; Feder, Gandal, Hamrick, Moore, et al., 2018).

In summary, despite being risky and poorly regulated, currency exchanges facilitated the diffusion and a widespread adoption of Bitcoin, by providing an online interface to conventional payment systems and thus fostering an efficient allocation of resources between demand and supply. They reduce the barriers to entry and technological costs; they also exploit economies of scale that reduce overall costs. It is curious that exchanges are so important for the Bitcoin ecosystem, whilst they are a Bitcoin superstructure that dramatically favors the dynamics of centralization and power concentration, where users accept to lose the direct control of their own coins, and essentially functioning through a design at the opposite of the one that constitutes and characterises the innovative core of the Bitcoin technology.

Chapter 2

The leaked Mt. Gox dataset: preprocessing, and an economic overview

This Chapter describes the Mt. Gox leaked dataset, the primary and most important source of information used to conduct the analyses. It contains the log of all the internal trades performed in the exchange platform in between 1 April 2011 and 30 November 2013.

First, I compare the deduplication methods and cleaning procedures I implemented to the ones used in the literature. In particular, I show that my procedure combines and comprises the former ones by improving upon them in some ways that I will clarify below. Then I exploit the preprocessed data to provide an economic overview of the Mt. Gox ecosystem, and how it evolved over time. Finally, I propose a model to estimate the fees a user would expect to pay given the official fee schedule published by the exchange platform.

Besides this model for the expected explicit transaction costs, the main novelty of this chapter consist in the provision of a public and preprocessed version of the leaked Mt. Gox dataset. In this sense, Section 2.1 is especially useful as it also provides a tool to understand both the meaning of the variables and the steps necessary to clean the dataset.

2.1 Dataset cleaning procedures

The first part of the polishing procedures is devoted to identify and remove the redundant duplicate rows contained in the leaked dataset. I follow the methodology described in the literature by Feder, Gandal, Hamrick, and Moore (2018), Gandal, Hamrick, et al. (2018) and Scaillet, Treccani, and Trevisan (2017). I discuss which files of the leaked dataset can be safely discarded without any loss of information and control for the presence of duplicate rows in the remaining files. Second, I further remove the rows with misreported entries. In some cases the correct values can be retrieved: thus, when possible, I chose to extrapolate the plausibly correct values. Finally, I conduct additional sanity checks on the quality of the data: I verify the correctness of the filtered data by comparison to external sources of information, that is, a dataset made public by Mt. Gox, and aggregated data from the website Bitcoincharts.com[1], a benchmark for many studies on cryptocurrency markets that collects anonymously information at the trade level for the main cryptocurrency exchange platforms.

2.1.1 Description of the leaked log files

The analysis is conducted over a total of 62 CSV monthly files that cover a time period ranging from April 2011 to November 2013. Once merged, they amount to 22,175,247 rows per 19 columns. Overall, they share a common structure: the data are reported as a sequence of trades, each identified by a *trade ID*, and each composed of two rows corresponding to a buy and a sell leg. Most importantly, each leg is associated to a *user ID*. A trade can be schematized as follows:

Date	Trade ID	User ID	Type	Bitcoins	Money	Currency	Fees
T	0	1	buy	Amount_{BTC}	Amount_{CUR}	CUR	F1
T	0	2	sell	Amount_{BTC}	Amount_{CUR}	CUR	F2

[1] http://api.bitcoincharts.com/v1/csv/

Along with the standard matching mechanism, from September 2011 on Mt. Gox allowed to trade bitcoins against other currencies. To facilitate the trading in illiquid markets, the platform introduced also a feature called multi-currency trading, allowing investors to match orders even if buy and sell sides were executed against different fiat currencies (at the cost of an additional fee). This matching mechanism, slightly more complex than the single currency trade, requires the administrator of the exchange platform to act as an intermediary between parties. These trades represent a marginal fraction of the whole log of trades.

The single leaked files largely overlap: several legs are duplicates and the data are not reported in a homogeneous format. Some patterns can anyway be identified. First, for each month it is possible to identify one primary file that contains all the relevant information, while the remaining ones are redundant subsets. Second, the whole body of files can be grouped in 3 macro sets, each sharing the same structure; the differences across them mainly concern the way multi-currency trades are transcribed, and the columns stored. Thus, files are divided and analyzed in three blocks, sharing a similar pattern:

1. files related to **April 2011**. Two different files are reported for April 2011, both composed of 15 columns: *Trade_Id, Date, User_Id, Japan, Type, Currency, Bitcoins, Money, Money_Rate, Money_JPY, Money_Fee, Money_Fee_Rate, Money_Fee_JPY, Bitcoin_Fee, Bitcoin_Fee_JPY*; multi-currency trades are not present.

2. files from **May 2011 to October 2012**: for every month there is only one .CSV file, each with the same columns of the April 2011 files. Multi-currency trades are implemented as follows:

User ID	Type	Bitcoins	Money	Currency
1	buy	A_{BTC}	A_{CUR1}	CUR1
TIBANNE_LIMITED_HK	sell	A_{BTC}	A_{CUR1}	CUR1
2	sell	A_{BTC}	A_{CUR2}	CUR2
TIBANNE_LIMITED_HK	buy	A_{BTC}	A_{CUR2}	CUR2

I will refer to this implementation method as the 'Tibanne' one;

3. files from **November 2012 to November 2013**. Two files per month are available, one denominated 'Coinlab' and one denominated 'mtgox_japan'. In addition, for February, March, and April 2013, also weekly datasets (with a similar nomenclature) are included. All of them are composed of 19 columns: the 15 aforementioned, plus *User, User_Id_Hash, User_Country, User_State*; the 'Coinlab' and weekly files are subsets of the original 'mtgox_japan' ones. The multi-currency trades are stored in the database as follows:

User ID	Type	Bitcoins	Money	Currency
3	buy	A_{BTC}	A_{CUR1}	CUR1
THK	sell	A_{BTC}	A_{CUR1}	CUR1
3	buy	A_{BTC}	A_{CUR1}	CUR1
THK	sell	A_{BTC}	A_{CUR1}	CUR1

I will refer to this implementation method as the 'THK' one. Note that the trades are misreported, as already noticed in the literature: the buy side of the trade is reported twice. The sell side of the 'THK' multi-currency trades is never reported.

4. The file related to **July 2012** represents an exception to this classification: it shares the same structure and properties of the second group, but the multi-currency trades follow the 'THK' scheme. Note that in all the files the 'Tibanne' and 'THK' methods are mutually exclusive.

The 19 columns can be interpreted in the following way:

- *Trade_Id*: identifier of the trade. Each pair of buy and sell legs is associated to a specific identifier; the identifiers are reported as a simple sequential increasing number until 19 June 2011, while from 26 June 2011 on, as suggested in Scaillet, Treccani, and Trevisan

(2017), they correspond to the concatenation of a POSIX timestamp and a microsecond timestamp[2];

- *Date*: date of execution of the trade. Time is reported in format YYYY-MM-DD hh:mm:ss;
- *Type*: value that determines whether the row represents a buy or a sell leg. The object of exchange is an amount of bitcoins; thus the buyer purchases bitcoins denominated in fiat currency, and viceversa for the seller;
- *User_Id*: parameter that identifies the user who performed the buy or sell action;
- *Japan*: the meaning of this variable is not clear. It can take two values: 'JP' and 'NJP'. It could be a parameter informative on the geographical origin of the trade execution, or it might indicate a special category of users (or trades);
- *Currency*: fiat currency used in the trade. In total, 17 different currencies are used: US Dollar, Euro, British Pound, Polish Zloty, Australian Dollar, Japanese Yen, Canadian Dollar, Swedish Krone, Swiss Franc, Russian Ruble, Chinese Yuan, New Zealand Dollar, Singapore Dollar, Hong Kong Dollar, Danish Krone, Norwegian Krone, Thai Bat.
- *Bitcoins*: amount of bitcoins exchanged;
- *Money*: quantity of fiat currency traded;
- *Money_JPY*: quantity of fiat currency traded, expressed in Japanese Yen;
- *Money_Rate*: exchange rate used to convert the value in the field 'Money' into Japanese Yen;

[2]On 19 June 2011, Mt. Gox's website went down for several days after a security breach. See `https://bit.ly/3dtzNFf`. This link and the following were all accessed on 14 October 2020.

- *Money_Fee* and *Bitcoin_Fee*: fees paid to perform the trade. Normally, fees were paid in fiat money by the seller and in bitcoins by the buyer (with a certain degree of flexibility on the choice of how to pay the commission fees[3]);
- *Money_Fee_JPY* and *Bitcoin_Fee_JPY*: fees paid to perform the trade, expressed in Japanese Yen;
- *Money_Fee_Rate*: exchange rate used to convert the value in the fields 'Money_Fee' and 'Bitcoin_Fee' into Japanese Yen;
- *User*: user identifier expressed in hexadecimal base;
- *User_Id_Hash*: hashed representation of the user identifier;
- *User_Country*: National geographic location of the user;
- *User_State*: Regional geographic location of the user.

2.1.2 Deduplication methods - comparisons with the literature

Members of the Mt. Gox community were among the first to explore the leaked dataset. Supposedly they wanted to prove misbehavior of the exchange in the events that led to its bankruptcy on 28 February 2014. The volunteers analyzed the structure of the dataset, tried to identify potential malicious users, and pointed out key issues to keep in mind[4]. I followed their example, and replicated the steps adopted by Feder, Gandal, Hamrick, and Moore (2018) and Gandal, Hamrick, et al. (2018).

The authors use two related methods to detect duplicates. The first one (method *Conservative*) detects rows as duplicates if the following entries are equal: *user id, timestamp, buy/sell action, amount in BTC, amount in Yen*. The other method detects rows as duplicates if the following entries are equal: *user id, timestamp, buy/sell action, amount in BTC*. The latter

[3]See https://bit.ly/34Wyb3h.

[4]See, for example: https://bit.ly/34WhQvo, https://bit.ly/2FtFtCJ, https://bit.ly/3iWtBXB, https://bit.ly/3dt0gD4, and https://bit.ly/3lQi8e9.

is more aggressive because it removes a higher number of rows, hence previous works refer to it as *Aggressive*.

However, these approaches also treat as duplicates unwanted legs, and do not take into account the likely presence of metaorders in the leaked dataset. Consider the case in which a user performs two exactly equivalent trades at the same moment, with the only difference that the complementary leg is executed by different trading partners: both deduplication methods mentioned above remove one of the two exactly equivalent legs. Thus, these methods reduce the dataset more than desirable.

To prevent this behavior, I slightly changed the deduplication *Aggressive* method, by adding the *trade id* value to the set of variables used to detect duplicates as in Scaillet, Treccani, and Trevisan, 2017 (method *TradeId*). As a result, rows are detected as duplicates if the following entries are equal: *trade id, user id, timestamp, buy/sell action, amount in BTC*. To compare the results, I also implemented another deduplication technique (method *Pairs*), based on the *Aggressive* method, but the legs of a trade are not treated independently: rows are considered as duplicates *only if* both legs are duplicates.

To clarify the differences among the different methods, a series of example trades and the resulting deduplications are shown in the following.

Original sample. Table 1 shows the original table. It also corresponds to the deduplication results of method **TradeId**, meaning that in that specific case the *TradeId* method does not remove any duplicate. Consider the following example: rows 937 and 939 show two equal legs, having the same values for *user id, timestamp, buy/sell action, amount in BTC, amount in Yen*; however, since the *trade id* marks them as distinct, they are not treated as duplicates.

Conservative. Here, instead, row 939 is considered a duplicate of row 937. To maintain the dataset coherent, both rows 939 *and* 938 are removed. Row 935 is not a duplicate of row 933 because of the difference in the value of 'Money_JPY'. Table 2 shows the result.

Aggressive. In every trade User 388 is seller at the same date and

Table 1: Original sample and result of *TradeId* deduplication method

	Trade_Id	Date	User_Id	Type	Bitcoins	Money_JPY
930	35837	11-04-04 14:23	2824	buy	10.0	586.89
931	35837	11-04-04 14:23	388	sell	10.0	586.89
932	35838	11-04-04 14:23	3111	buy	10.0	578.42
933	35838	11-04-04 14:23	388	sell	10.0	578.42
934	35839	11-04-04 14:23	2824	buy	10.0	570.20
935	35839	11-04-04 14:23	388	sell	10.0	570.20
936	35840	11-04-04 14:23	3111	buy	10.0	570.00
937	35840	11-04-04 14:23	388	sell	10.0	570.00
938	35841	11-04-04 14:23	1000	buy	10.0	570.00
939	35841	11-04-04 14:23	388	sell	10.0	570.00

Table 2: Result of *Conservative* deduplication technique

	Trade_Id	Date	User_Id	Type	Bitcoins	Money_JPY
930	35837	11-04-04 14:23	2824	buy	10.0	586.89
931	35837	11-04-04 14:23	388	sell	10.0	586.89
932	35838	11-04-04 14:23	3111	buy	10.0	578.42
933	35838	11-04-04 14:23	388	sell	10.0	578.42
934	35839	11-04-04 14:23	2824	buy	10.0	570.20
935	35839	11-04-04 14:23	388	sell	10.0	570.20
936	35840	11-04-04 14:23	3111	buy	10.0	570.00
937	35840	11-04-04 14:23	388	sell	10.0	570.00

Table 3: Result of *Aggressive* deduplication method

	Trade_Id	Date	User_Id	Type	Bitcoins	Money_JPY
930	35837	11-04-04 14:23	2824	buy	10.0	586.89
931	35837	11-04-04 14:23	388	sell	10.0	586.89

quantity. Thus, independently on the partner, all are considered as duplicates of the trade at rows 930, 931. Results are shown in Table 3.

Pairs. Here, instead, we remove only trades where *both* legs of a trade are duplicates according to the criterion *user id, timestamp, buy/sell action, amount in BTC*. Note that Trade Id is *not* considered to detect duplicates. As depicted in Table 4, pairs 934, 935 and 936, 937 are removed, while pair 938, 939 is kept, given the presence of a different user w.r.t. previous trades in the *buy* side.

Table 4: Result of *Pairs* deduplication method

	Trade_Id	Date	User_Id	Type	Bitcoins	Money_JPY
930	35837	11-04-04 14:23	2824	buy	10.0	586.89
931	35837	11-04-04 14:23	388	sell	10.0	586.89
932	35838	11-04-04 14:23	3111	buy	10.0	578.42
933	35838	11-04-04 14:23	388	sell	10.0	578.42
938	35841	11-04-04 14:23	1000	buy	10.0	570.00
939	35841	11-04-04 14:23	388	sell	10.0	570.00

I follow the approach introduced in Scaillet, Treccani, and Trevisan, 2017 and conduct the analyses using the *TradeId* deduplication method. As a robustness check, I also repeat them using the *Pairs* method; results are coherent.

2.1.3 Identification of the redundant files and deduplication

I verify which files are redundant. The dataset includes two different versions for April 2011: the two differ by 41 rows, and only with respect to the field 'User_Id': all these legs are reported as executed by user 634 in the first, while this identifier was hidden (User_Id = 'DELETED') in the latter. All the 'Coinlab' files from November 2012 to November 2013 contain only duplicated rows when compared to the corresponding monthly 'mtgox_japan' ones. Instead, a set of rows in the weekly 'mtgox_japan' files apparently result as non-duplicates, all of them being performed by user 634. However, upon closer inspection, it emerges that also these rows are duplicated, if taking into account posthumous corrections to the fields 'Money' and 'User_Id': in each of these instances, the user identifier is systematically changed from 698630 to 634, and the values in the 'Money' field, seemingly random, are corrected to market values. Consistently with the literature, I assume the rows with user ID 698630 to be the original ones. These trades were likely not supported by a real transfer of funds from an account to another, althought virtually executed: thus, as noted by Gandal, Hamrick, et al., 2018 (p. 89) "no legitimate Mt. Gox customer received the currency Markus supposedly paid to acquire these claimed coins"[5]. However, the parties involved in these trades with user 698630 were supposedly not aware of the underlying illicit activity; for them, the trade had normally occurred[6]. Thus, for the purpose of my analysis, it is correct to consider these trades as if they had been executed at market prices.

Once merged, after discarding the redundant files, the dataset is composed of 16,748,681 rows and 19 columns. Using the two deduplication techniques described before (*Pairs* and *TradeId*), I control for the presence of further duplicate trades in the remaining files. Note that, in this step, the latter methodology does not reduce further the dataset.

[5]For further details, see also the online supplementary material `https://bit.ly/3jZDvcp` of Gandal, Hamrick, et al., 2018.

[6]One can verify that these trades had a market impact, see e.g. in bitcoincharts.com (`https://bit.ly/39vT9tx`): the large price variations are due to user 698630's activity.

2.2 Dataset sanity checks

Figure 4: Exchange rate of a sample of USD trades in the merged dataset before the polishing procedures.

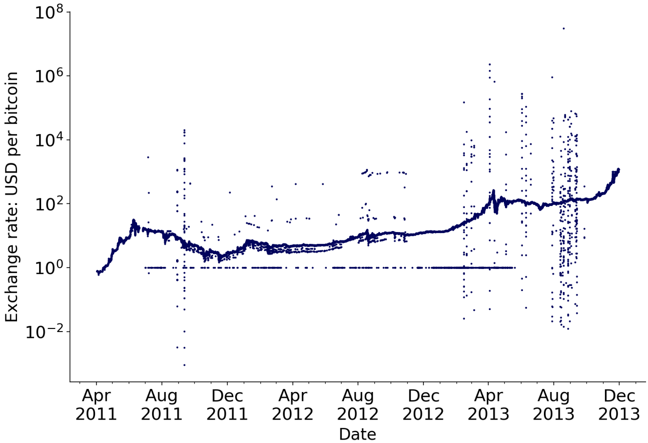

Figure 4 shows the computed exchange rates for a sample of trades, extracted from the merged file, executed in USD. Many values appear to be not related to market values. Given the relevance of this variable for the analysis, and to guarantee high quality standards for the data used, I perform additional sanity checks built on approaches used in the literature to further polish the dataset. The reported values refer to the *TradeId* deduplcation method:

- Last day trades are removed, due to the presence of several inconsistencies;
- 115,275 self-trades are removed[7];
- rows with user ID 'TIBANNE_LIMITED_HK' and 'THK' are removed (they do not represent trades; rather, these are redundant

[7]The figures reported are comparable with those found in Scaillet, Treccani, and Trevisan, 2017. The difference is due to the fact that, in their analysis, they discard all non USD trades.

Figure 5: Exchange rate of trades executed by user 698630 in March (left) and June (right) 2013.

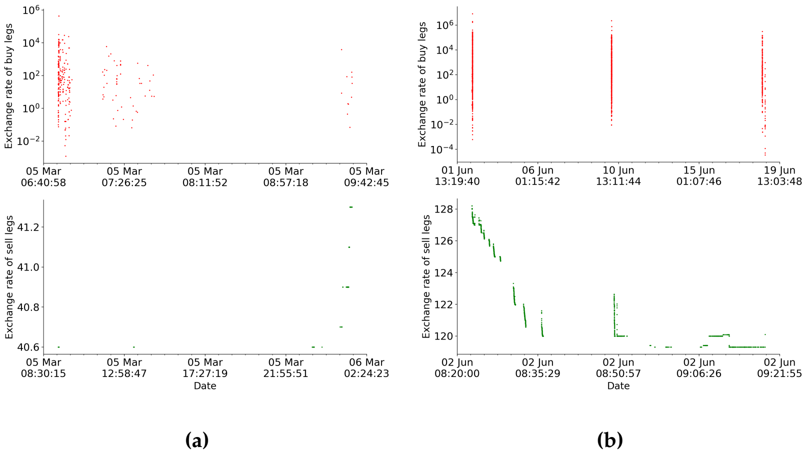

Notes: the choice of the months is illustrative. Buy trades (upper panels) have random exchange rates, while exchange rates of the sell trades (lower panels) are at the market prices.

rows representing the mediation of the platform among parties for multi-currency trades);

- Around 136,000 trades (92,000 in USD) are systematically misreported. All these are multi-currencies trades of the 'Tibanne' type. The 'Money' entry is the same in the primary leg (the buy side) and in the secondary leg (the sell side), thus the order of magnitude of the error is given by the exchange rate between the two currencies. I do not remove these legs as the correct data can be retrieved;

- 336,813 rows, all secondary legs belonging to multi-currency trades based on the 'THK' method, are missing (the dataset reports twice the primary leg). It is not possible to recover the data for these missing legs;

- 31 trades with 'DELETED' user identifier or the field 'Bitcoins' = 0 are removed;

- I verify that the dates are expressed in the same time zone (UTC);

- following Gandal, Hamrick, et al., 2018 and anonymous contributions from the Mt. Gox users community, I identify two key users:

  - The literature refers to the account 698630 as **Markus**, and suspects pertain to the legality of its actions. It never paid fees, and all the trades where it appears as a buyer have seemingly random values in the 'Money' field (supposedly because the virtual trades were not backed by real transactions: as a consequence, the trading log mechanism would interpret incorrectly the void entry and would fill it by simply copying and pasting the last 'Money' value previously transcribed in the log). Figure 5 shows its trade pattern, distinguishing between buy and sell legs. Notice that this account is strictly linked to the one with user ID = 634, as described before; user 634, active only in the first months of the dataset, seems to perform licit trades.

- The nickname **Willy** was given to 49 different accounts, bots sharing the same trade pattern and controlled by the same entity[8]: their 'User_Country' field is always marked as '??'; they only perform buy actions; these accounts activate one by one, once the previous has spent a definite round amount of dollars (usually $2.5 millions).

  While Markus is active from February 2013 to September 2013, Willy's bots become active a few hours after the last Markus' trade, suggesting they are plausibly linked.

It is likely that the same person owned all these accounts (634, 698630, and the 49 bots).

- The field 'User_Id' for Markus trades performed by the account '698630' is changed to 635 (after controlling that no user has user identifier = 635), so that the accounts 698630 and 634 are sequentially linked, but they can be treated as a single user; the field 'User_Id' for Willy's bots is changed to 1000000, a number which is not owned by any other user.

- I add a column that maps all the user identifiers in the dataset to a consecutive sequence of integers, preserving the order but not the numbers. Throughout the analysis I will use this mapping to guarantee user's anonymity;

- I add a column to retrieve multi-currency trades, where the entries take the following values:

  0. Standard trades;
  1. Tibanne multi-currency trades;
  2. THK multi-currency trades.

[8] https://bit.ly/34VCsUx.

2.3 Comparisons with other sources of information

2.3.1 The public Mt. Gox dataset

To further ensure the validity of the data, I compare the leaked log to a dataset made public by Mt. Gox and used by Scaillet, Treccani, and Trevisan (2017) in their work. This dataset contains the history of trades executed in Mt. Gox, from its birth on 17 July 2010 to 17 December 2013, for a total of 8,605,998 trades. The main difference with the leaked dataset is that trades are not reported as a couple of legs, but as a single trade, thus they do not contain information on the users who performed the buy and sell actions; however, the trade identifiers follow the same scheme used in the leaked dataset, allowing to merge the two datasets.

Besides general fields already included in the leaked log (such as the trade identifier, the fiat money used, the volume traded, the exchange rate, the date), the public dataset contains additional information on the trades, specifying the typology of the order (possible values are 'limit'; 'market'; 'limit,mixed_currency'; 'market,mixed_currency') and which party initiated the trade ('bid' or 'ask'). While there is no way to fully verify the correctness of these datasets, the comparison with a second source of information provides a solid robustness check.

I first verify that for each trade *the amount of bitcoins traded correspond exactly* in the two datasets. The values related to the money traded, instead, show some differences for around 310,000 legs. Data in the public dataset are reported as exchange rates, thus the following comparisons concern exchange rates. First, I find out that trades executed before 12 September 2013 in SEK and JPY are misreported in the public dataset by a factor of 10^2, explaining the divergences for around 32,000 legs. Second, around 135,000 multi-currency 'Tibanne' trades are wrongly transcribed in the leaked log, as the two legs report the same fiat amount; the comparison with the prices in the public dataset allows to correct them. Around 55,000 trades are transactions that involve small amounts of bitcoins (less than 10^{-3} BTCs), and in 8,000 the difference between exchange rates is smaller than 1% - suggesting for both those groups that the differences are due to different roundings. Among the remaining 'unexplained' legs (around 80,000), the vast majority is related to trades executed by user 635 (around 23,000 trades) and few other users: the first four users alone account for slightly less than 50,000 legs. Upon closer inspection, these errors seem to be caused by the same mechanism explained above for the Markus buy trades: the 'Money' field, when the value is void, is apparently filled with the entry reported in the previous trade. E.g., an account performed more than 2,500 nearly consecutive trades in day 2011-09-11, the vast majority of which with random 'Money' values.

Figure 6 shows the results for the USD market after the cleaning procedures. Figure 7 shows the exchange rates for other markets (the choice is arbitrary and for illustrative purposes): I plot the monthly BTCtoEUR exchange rate for March 2012 on the left, and the BTCtoGBP rate for April 2013 on the right.

Figure 6: Exchange rate of a sample of USD trades in the merged dataset after the polishing procedures and the comparison with the public dataset.

Figure 7: Exchange rate of a sample of trades in EUR (March 2012) and in GBP (April 2013).

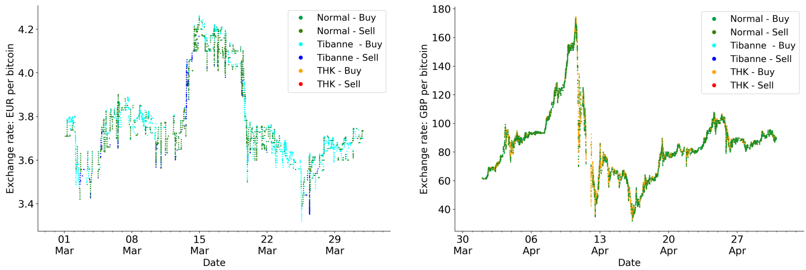

(a) **(b)**

Notes: the choice of the reported months is purely illustrative. The colors represent different types of trades: Tibane multi-currency trades (blue), THK mutli-currency trades (orange), normal trades (green).

2.3.2 Volume comparison with Bitcoincharts data

To get a form of external validation, I compare the daily USD volumes in the leaked dataset with the data made public by Bitcoincharts.com. This website provides reliable data and is a benchmark for many studies on the Bitcoin exchange platforms. The following plots compare USD denominated volumes aggregated on a daily basis. They represent the difference of volumes between the leaked and the bitcoincharts dataset, normalized on the former. The lines represent the daily differences and their 15 days centered moving average. Figure 8a shows the volume difference *after* the deduplication procedure using the method 'Trade Id', but *before* the cleaning procedure reported in Section 2.2.

The volumes differ only in the months were the 'THK' multi-currency method is implemented. As argued before, the secondary legs of such trades are missing. It is likely that some of them were in USD, plausibly explaining the missing volumes. Thus, Figure 8a can be interpreted *as if* no THK secondary leg were in USD; that is, as a 'lower bound' for USD-denominated volumes traded. Viceversa, Figure 8b plots the volume differences one would observe if all the secondary legs - of the THK trades whose first leg is not in USD - were in USD. This can be considered as an 'upper bound' to USD volumes traded. The plots suggest that this intuition is correct and that the difference between the Mt. Gox dataset and the Bitcoincharts data is due only to the missing multi-currency THK secondary legs. Thus this plot also quantifies the cost in terms of information loss due to the misreported sell side for such trades.

Figure 8: Volumes comparison with Bitcoincharts data.

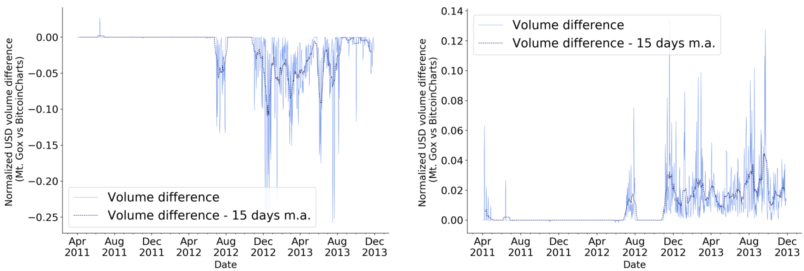

(a) **(b)**

Notes: the comparison is *before* performing the dataset sanity checks, using the deduplication method 'TradeID'. Panel (a) reports the differences between the daily volumes in the leaked dataset and the Bitcoincharts data; Panel (b) reports the 'upper bound' of the USD traded volumes that would be observed if all the missing secondary THK legs were in USD.

Finally, I remove the primary legs of the 'THK' multy-currency trades, to have a coherent dataset where all the reported trades have both the buy and the sell side. In Figure 9 I compare how the volume differences change when including (left) and excluding (right) the 'THK' trades, *after* implementing the sanity checks in Section 2.2. Note that most of the 'THK' removed trades were executed after April 2013.

2.3.3 Comparison with other results in the literature

Last, I identify the trading activity of the users Markus and Willy, and I report the aggregated amount of bitcoins and fiat currency that they

Figure 9: Volumes comparison with Bitcoincharts data (II).

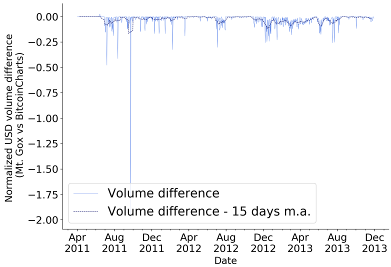

(a)

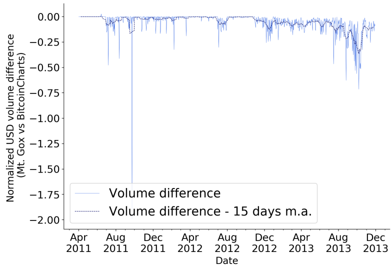

(b)

Notes: the comparison is *after* performing the dataset sanity checks, using the deduplication method 'TradeID'. Panel (a) reports the difference when the sell sides of the THK trades are included; Panel (b) is without the THK trades. The differences concern primarily the months after April 2013.

exchanged according to my deduplication approaches. I compare these figures with the findings in Gandal, Hamrick, et al., 2018. The results slightly change depending on the deduplication method used, but are consistent with those found previously in the literature (see Table 5).

To conclude, this section proves that the results of my deduplication methods are consistent with those implemented in previous analyses. Several checks are performed to ensure the quality of the dataset, and show that my procedure provides a high quality dataset which is consistent with external sources of information largely accepted in the literature. Further, I merge the leaked dataset with public information on the Mt. Gox trades, and I provide it in order to ensure reproducibility.

The final dataset I use is deduplicated using the *TradeId* methodology, and all the primary legs of the THK trades, whose secondary leg is missing, are excluded. It is composed of 14,875,192 rows, for a total of 7,437,596 trades. The vast majority of them are in USD (87.9%) and in EUR (7.7%), and are executed by 125,755 users.

Table 5: Aggregate information for the users *Markus* and *Willy*

Pairs without THK:	*Pairs with THK*:
Bitcoins bought by Markus: 302,928.392	Bitcoins bought by Markus: 330,994.764
Bitcoins sold by Markus: 35,535.441	Bitcoins sold by Markus: 35,535.441
Dollars spent by Willy: 100,739,634.260	Dollars spent by Willy: 108,247,263.292
Bitcoins bought by Willy: 237,689.192	Bitcoins bought by Willy: 261,437.290
TradeId without THK:	*TradeId with THK*:
Bitcoins bought by Markus: 306,339.336	Bitcoins bought by Markus: 335,897.702
Bitcoins sold by Markus: 35,867.176	Bitcoins sold by Markus: 35,867.176
Dollars spent by Willy: 102,710,284.814	Dollars spent by Willy: 110,376,397.492
Bitcoins bought by Willy: 242,794.122	Bitcoins bought by Willy: 266,984.786[a]

[a] *If the 1,147.948 bitcoins traded by Willy in the last day (removed from the dataset) are accounted for, I obtain exactly the amount found by Gandal et al.: 268,132.734 bitcoins.*

2.4 The Mt. Gox ecosystem

2.4.1 An economic overview

These paragraphs provide descriptive and qualitative findings on the evolution in time of some interesting economic indicators of the Mt. Gox ecosystem. Figure 10 shows the geographical distribution of the users that performed at least one trade from November 2012 on (as this information is not available from trades executed previously). The most represented country is US; Australia, China, Russia, and European countries also play a prominent role. Figures 11 and 12 show how the average trade magnitude evolved in time. In Panel (11a) I plot the daily mean amount of bitcoins per trade, in Panel (11b) the daily mean amount of dollars exchanged per trade. The first decreases over time, as one would intuitively expect given the large positive bitcoin price variation. It is less straightforward, instead, to predict a priori the dynamics of the latter: Panel (b) shows that the mean amount of dollars exchanged per trade increased constantly in time. Figure 12 adds further information by plotting their respective quartiles.

Figure 10: Number of active users per country.

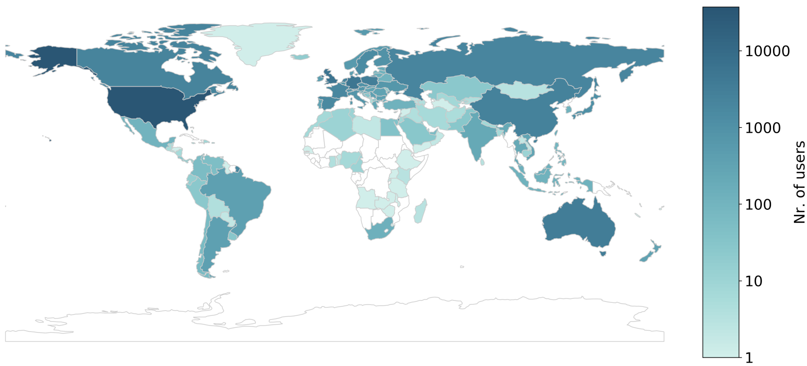

Figure 11: Average traded volume, evolution in time.

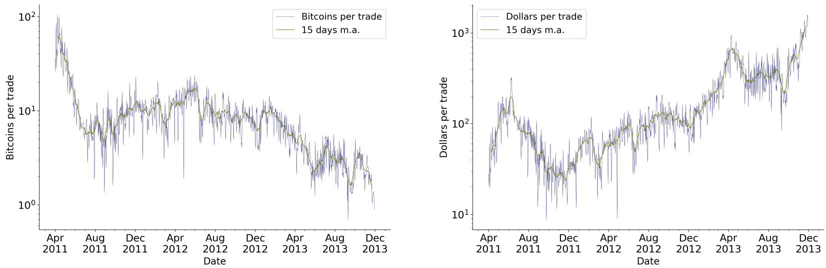

(a) **(b)**

Notes: Panel (a) shows the daily evolution of the mean amount of bitcoins per trade, while Panel (b) shows the mean amount of Dollars per trade (considering only trades denominated in USD).

Figure 12: Average traded volume, evolution in time - quartiles.

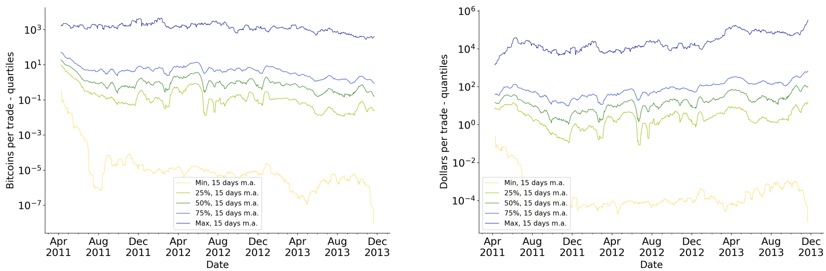

(a) **(b)**

Notes: quartiles for (a) the amount of bitcoins per trade, and (b) the amount of Dollars per trade (considering only trades denominated in USD).

Figure 13: Heterogeneous behaviors and fat-tailed distributions: rank-frequency plot for (a) the number of trades (Buy, Sell, Total) and (b) the days of activity.

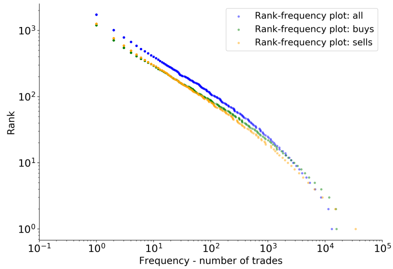

(a)

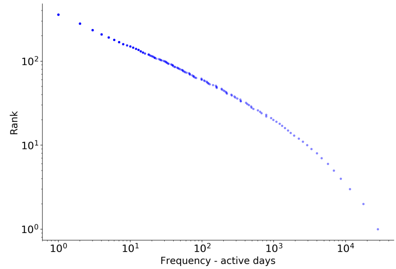

(b)

Figure 13 provides some insight on the users' heterogeneity; following the procedure described in Newman (2005), I constructed the rank frequency plot in Panel (13a) for the trades executed by each user, and in Panel (13b) for the number of active days of each user. I do not report the full analysis, but similar results hold for the distributions of Bitcoins bought and sold by each user, as well as for the fiat money that they bought and sold.

Finally, Figure 14 represents the 'daily activity line' for the first 1.000 users who traded inside the Mt. Gox platform: for each row representing a user U, column values (that represent days) are white if U was active during that day, and black if not active. The choice of the sample is illustrative.

Figure 14: Daily activity map for the first 1000 users. White: active, black: not active

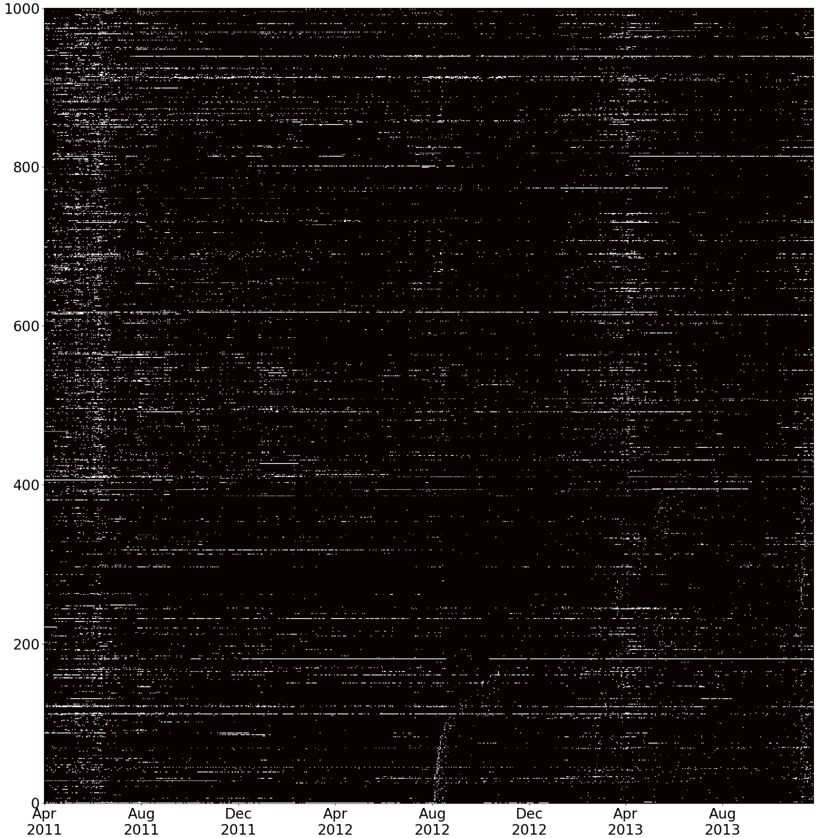

2.5 Fees Scheme

An essential feature of the dataset is that it includes the explicit transaction costs associated to the legs of each trade. Such additional costs are relevant for the purpose of the analysis in that they affect the profitabitlity of a trade, and consequently it is necessary to take them into account also to identify potentially profitable arbitrage actions, which is the primary purpose of the next chapters.

Due to the inconsistencies found in the leaked dataset and discussed above, and to overcome the risk that they might affect also the data regarding the fees paid by the users, in this section I provide and discuss a model that calculates the fees a user would theoretically pay given the fee schedule published by Mt. Gox. Indeed, according to their official posted scheme, the magnitude of the fees depends on the volume traded by the user in the last month — i.e., most active users in the recent past pay less.

Table 6: Mt. Gox posted fee schedule.

Volume (BTCs)	Fees	Volume (BTCs)	Fees
0 to $<$ 100	0.60%	10000 to $<$ 25000	0.30%
100 to $<$ 200	0.55%	25000 to $<$ 50000	0.29%
200 to $<$ 500	0.53%	50000 to $<$ 100000	0.28%
500 to $<$ 1000	0.50%	100000 to $<$ 250000	0.27%
1000 to $<$ 2000	0.46%	250000 to $<$ 500000	0.26%
2000 to $<$ 5000	0.43%	$>$ 500000	0.25%
5000 to $<$ 10000	0.40%	-	-

Notes: discounts are based on the user's trading volume over the last 720 hours.

Table 6 shows the fee schedule published by Mt. Gox; I could confirm that the scheme was valid at least from 16 October 2011 until mid-February 2013[9]. Discounts were based on the volume of bitcoins traded by the user over the last 720 hours[10] (that is, 30 days).

[9] https://bit.ly/374zXBW, https://bit.ly/2GYVcdF, and https://bit.ly/2IrPRMp

[10] https://bit.ly/2SWwJIx

Figure 15: Empirical fee schedule

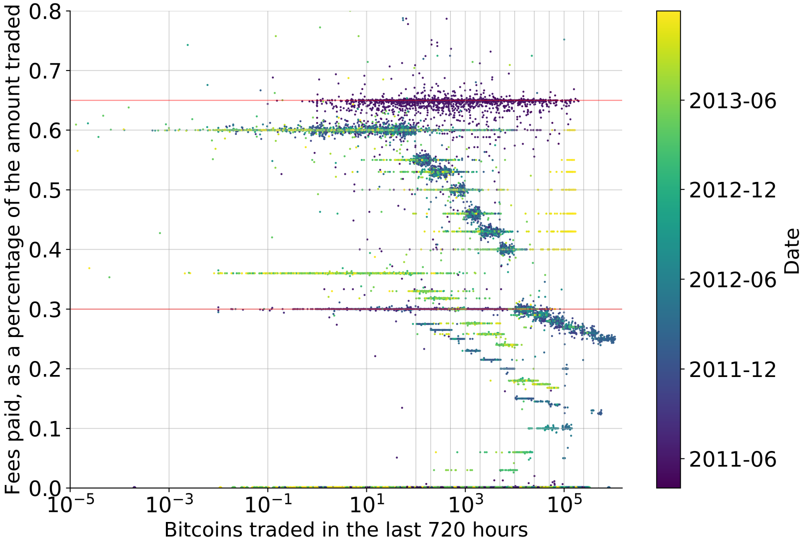

Notes: Fees paid per leg, expressed as the percentage of the value of the trade (y-axis). According to Mt. Gox's schedule, fees depend on the amount of bitcoins traded by the user in the last 720 hours (on the x-axis). Color indicates the time of the trade. The lines of the grid help to graphically delimit the volume bands corresponding to different discount brackets. Two horizontal red lines help to identify the two time windows when fees were fixed: 0.65% from 1 April 2011 to 23 June 2011, and 0.3% from 24 June 2011 to 18 August 2011. Many legs (more than one million) pay no fees, and a limited number of legs pays fees as low as 0.1 percent or less. In addition, in some particular circumstances (e.g., Easter and Christmas holidays in 2011 and 2012), fees were halved, thus explaining some of the values that do not correspond to the official fee schedule. The plot shows a random sample of $N = 100,000$ legs. Note that a small fraction of trades (around 15,000 over the total of 7.5 million trades) have fees equal or above 1% and are likely to correspond to misreported data. I explicitly decided to plot only the legs between 0% and 0.8%.

I first compare the posted and the real fee schedules. In the Mt. Gox dataset, transaction costs are reported in two entries: by default bitcoin buyers were charged with fees in bitcoin, while sellers with fees in fiat money, but they could also choose how to pay them[11]. Thus, for each leg of every trade, I compute the actual fees paid as

$$\text{Fee} = \frac{\text{BitcoinFee}}{\text{Bitcoins}} + \frac{\text{MoneyFee}}{\text{Money}}, \tag{2.1}$$

where 'BitcoinFee' represents the fees paid on the amount of bitcoins traded, while 'MoneyFee' represents the fees paid on the amount of fiat money traded.

Figure 15 shows a sample ($N = 100,000$) of the empirical fees, focusing on the relationship between the actual transaction costs and the past volume traded; each dot represents the fees paid on a leg of a trade. By comparing the posted and the actual schedules, one can see that most of the data points fall into the expected volume bands, although deviations exist: first, from 0.40% to 0.20%, many points follow a pattern that cannot be explained by the posted schedule; second, a non-negligible number of dots falls below the threshold of 0.20%, suggesting the existence of privileged users, and a subset of legs is completely exempted from any kind of fee: Figure 16 reports additional information on them.

For this reason, instead of reverse-engineering the posted schedule, I take an empirical approach and fit a simple model that predicts the fees a user would have to pay given his trading history. The fee model is specified as:

$$\begin{aligned} \text{Fee}_i = \beta_0 &+ \beta_1 \cdot \text{LogVol}_i + \beta_2 \cdot \text{VolSmall}_i + \beta_3 \cdot \text{VolBig}_i \\ &+ \beta_4 \cdot \text{LogVol}_i \cdot \text{VolSmall}_i + \beta_5 \cdot \text{LogVol}_i \cdot \text{VolBig}_i \\ &+ \beta_6 \cdot \text{T}_{0i} + \beta_7 \cdot \text{T}_{1i} + \beta_8 \cdot \text{T}_{\text{holid}i} + \epsilon_i. \end{aligned} \tag{2.2}$$

[11] https://bit.ly/34Wyb3h

The independent variables have the following meaning:

- LogVol, the natural logarithm of the volume traded in the last 720 hours by the user who submitted the leg associated to the fee;
- $\mathrm{VolSmall}$, a dummy variable equal to 1 if the volume traded in the last 720 hours is between 100 and 10,000 bitcoins;
- VolBig, a dummy variable equal to 1 if the volume traded in the last 720 hours exceeds 10,000 bitcoins. As it can be seen from Figure 15, discount factors follow a linear trend with different slopes below and above the 10,000 bitcoins threshold. This is the reason why I introduced this dummy variable and the previous one, as well as their interaction terms with the LogVol variable;
- T_0, a dummy variable for trades executed between 1 April 2011 and 23 June 2011;
- T_1, a dummy variable for trades executed between 24 June 2011 and 18 August 2011;
- $\mathrm{T}_{\mathrm{holid}}$, a dummy variable for trades executed on 'special days': from 26 December 2011 to 1 January 2012, from 2 to 7 April 2012; on 9 and 10 November 2012[12].

Table 7 reports the estimated coefficients and goodness-of-fit indicators. In each specification, the constant term approximates the non-discounted official fee of 0.6%, and the response variable is negatively correlated with an increase of the volume traded in the past 720 hours; as expected, when included in the model, T_0 and T_1 respectively increase and decrease the constant term by a factor of around 0.16%, while $\mathrm{T}_{\mathrm{holid}}$has an even stronger negative effect (around -0.19%). Finally, again in accordance with the expectations, both β_4 and β_5 are negative; the first one is bigger in absolute terms, thus indicating a steeper variation of discounts given the same variation in volume (as can be seen in Figure 15).

[12]respectively, Christmas holidays in 2011, Easter holidays in 2012 and first Bitcoin Friday Sale day.
`https://bit.ly/37317sY`

Table 7: Fee model for non-zero fees, coefficients fitted with OLS

Dependent variable:	Fee as a percentage of the amount traded				
	Specification				
	(1)	(2)	(3)	(4)	(5)
Intercept	0.638	0.640	0.557	0.559	0.561
	(0.0001)	(0.0001)	(0.0000)	(0.0001)	(0.0001)
LogVol	-0.030	-0.030		-0.001	-0.001
	(0.0000)	(0.0000)		(0.0000)	(0.0000)
VolSmall			-0.107	0.150	0.152
			(0.0001)	(0.0002)	(0.0002)
VolBig			-0.284	-0.201	-0.212
			(0.0001)	(0.0005)	(0.0004)
LogVol $*$ VolSmall				-0.037	-0.037
				(0.0000)	(0.0000)
LogVol $*$ VolBig				-0.007	-0.006
				(0.0001)	(0.0001)
T_0		0.164			0.158
		(0.0002)			(0.0001)
T_1		-0.162			-0.170
		(0.0001)			(0.0001)
T_{holid}		-0.187			-0.191
		(0.0002)			(0.0002)
R^2	0.516	0.610	0.533	0.598	0.694
Obs.	13083547	13083547	13083547	13083547	13083547

Notes: All variables are significant at the 0.1% level. This is due to the high number of observations; however, I emphasize that these results are intended not so much to find significant effects as to predict fees. Observations consist of legs whose fees are positive and their value is below 1%.

I chose a logarithmic model because the break points for the fee schedule in Table 6 follow a logarithmic trend, and thus it was straightforward to consider the estimation on logarithmic volumes. To strengthen the results, I explored an alternative model using linear volumes; I report the results in Table 8.

The comparison of the two models shows that the overall pattern is similar, but in the logarithmic models the coefficients can be interpreted with less decimal digits, and the parameters' variations of signs are less frequent; in addition to that, in linear models (1) to (3) the intercept is smaller than expected (0.45%) and the interaction terms β_4 and β_5 are close to zero, in contrast with what one would expect. Regarding the explanatory power, the logarithmic model (1) outperforms the same model with linear volume (R^2 = 0.510 versus R^2 = 0.097), as well as the estimations (2) and (3). Model (4) and (5) have comparable explanatory power, but also in that case the logarithmic model has higher R^2.

Even though not directly comparable, in the work of Kim (2017) on the cost advantage of Bitcoin over cross-border ATM transactions the models for Bitcoin cross-border transaction costs achieve R-squared values in the order of 0.5 for the time period 2014-2015. Thus, it seems that cryptocurrency fees can be explained with linear regressions at this order of magnitude.

Finally, I propose a logit model (Table 9) to estimate the probability that a leg pay zero fees given user-specific and time-related variables: here the binary dependent variable is 1 if the trader paid some fee, and 0 otherwise. The model is specified as follows:

$$\begin{aligned} log\left(\frac{\mathrm{Fee}_i}{1-\mathrm{Fee}_i}\right) = \beta_0 + \beta_1 \cdot \mathrm{LogVol}_i + \beta_2 \cdot \mathrm{Bitcoins}_i + \beta_3 \cdot \mathrm{Date}_i + \\ +\beta_4 \cdot \mathrm{AnomalousDays}_i + \beta_5 \cdot \mathrm{EarlyAdopters}_i + \\ +\beta_6 \cdot \mathrm{AnomalousUsers}_i + \beta_7 \cdot \mathrm{Matchers}_i + \\ +\beta_8 \cdot \mathrm{Markus}_i + \beta_9 \cdot \mathrm{Willy}_i + \epsilon_i \end{aligned} \tag{2.3}$$

and the independent variables have the following meaning:

- $\mathrm{Bitcoins}$: the amount of bitcoins traded;
- Date: date of the trade execution;
- $\mathrm{EarlyAdopters}$: dummy variable equal to 1 for the first $16,000$ user IDs in sequential order;
- $\mathrm{AnomalousDays}$: dummy variable equal to 1 for the days with an anomalous presence of zero fees trades;

Table 8: Fee model for non-zero fees, coefficients fitted with OLS (linear alternative, not used)

Dependent variable:	Fee as a percentage of the amount traded				
	Specification				
	(1)	(2)	(3)	(4)	(5)
Intercept	0.454	0.456	0.469	0.559	0.560
	(0.0000)	(0.0000)	(0.0000)	(0.0001)	(0.0001)
LinVol	-0.000	-0.000		-0.000	-0.000
	(0.0000)	(0.0000)		(0.0000)	(0.0000)
VolSmall				-0.071	-0.069
				(0.0001)	(0.0001)
VolBig				-0.280	-0.281
				(0.0001)	(0.0001)
LinVol $*$ VolSmall			-0.000	0.000	0.000
			(0.0000)	(0.0000)	(0.0000)
LinVol $*$ VolBig			-0.000	0.000	0.000
			(0.0000)	(0.0000)	(0.0000)
T_0		0.169			0.158
		(0.0002)			(0.0001)
T_1		-0.149			-0.171
		(0.0002)			(0.0001)
T_{holid}		-0.202			-0.192
		(0.0003)			(0.0002)
R^2	0.103	0.197	0.144	0.585	0.682
Obs.	13083547	13083547	13083547	13083547	13083547

Notes: All variables are significant at 0.1% level. This is due to the high number of observations; however, I emphasize that these results are intended not so much to find significant effects as to predict fees. Observations consist of legs whose fees are positive and their value is below 1%.

Figure 16: Patterns of waived fees.

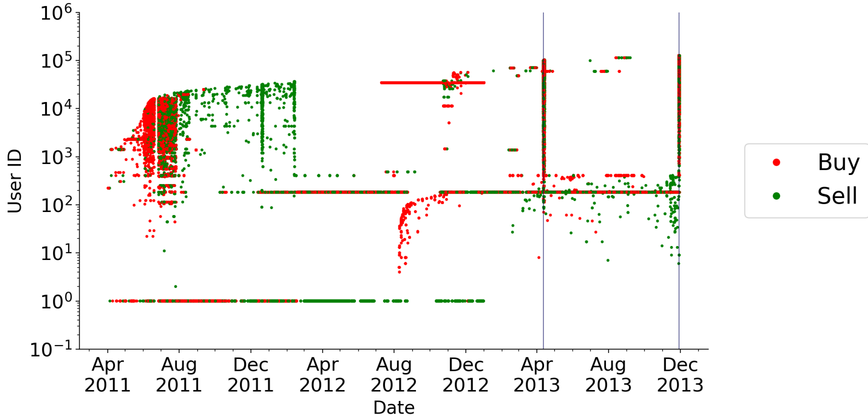

(a)

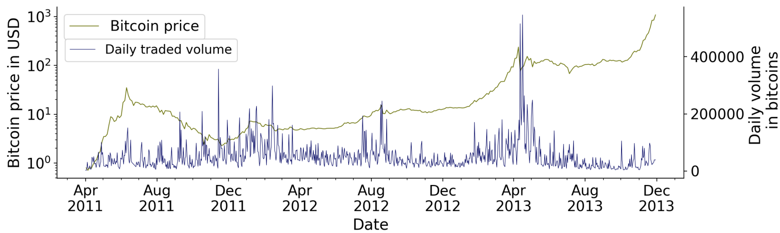

(b)

Notes: Panel (a) reports the distribution of the legs whose associated fees are zero, as a function of time and user ID (the latter on a logarithmic scale, to focus on low IDs). Buy (red) and sell (green) orders differ by color. Many interesting patterns emerge: first, two users with low ID did not pay fees over extended time periods; moreover, they account for $\sim 1,000,000$ zero-fee legs over the total of $\sim 1,700,000$ zero-fee legs. Second, during some days (especially on 19, 20, and 21 December 2011; 12, 13, and 14 April 2013; 28 and 29 November 2013) there were anomalous increases in trade legs with zero fees. Possible explanations include special events, such as a temporary downtime of Mt. Gox on 11 April 2013, and the exchange rate surpassing 1,000 \$/BTC for the first time on 27 November 2013. In both cases, the number of zero-fee trade legs increases shortly afterwards. Panel (b) depicts the daily bitcoin volume traded and the exchange rate in USD (the latter on a logarithmic scale). Observe from both panels that many users with low IDs seem to have *coordinately* executed buy orders from August to around November 2012, and then exclusively sell orders in the the days preceding the Bitcoin price peak. The plot shows a random sample of $N = 200,000$ legs.

Table 9: Fee model for zero fees, coefficients fitted with the logistic regression

Dependent variable:	Dichotomous variable accounting for the presence of fees				
	Specification				
	(1)	(2)	(3)	(4)	(5)
Intercept	5.104 (0.0033)	4.429 (0.0046)	1.670 (0.0019)	3.246 (0.0034)	3.257 (0.0073)
LogVol	-0.378 (0.0003)	-0.314 (0.0004)			-0.050 (0.0007)
Bitcoins		0.011 (0.0001)			0.005 (0.0001)
Date		-0.001 (0.0000)	0.001 (0.0000)		0.002 (0.0000)
AnomalousDays			-3.402 (0.0034)		-6.459 (0.0060)
EarlyAdopters				-0.897 (0.0037)	-1.575 (0.0055)
AnomalousUsers				-5.556 (0.0044)	-6.941 (0.0065)
Matchers				0.929 (0.0052)	0.799 (0.0062)
Markus				-1.434 (0.0156)	-1.468 (0.0159)
Willy				-0.418 (0.0189)	-0.675 (0.0306)
$pseudo - R^2$	0.157	0.228	0.109	0.472	0.676
Obs.	14875192	14875192	14875192	14875192	14875192

Notes: All variables are significant at 0.1% level. This is due to the high number of observations; however, I emphasize that these results are intended not so much to find significant effects as to predict fees.

- $\mathrm{AnomalousUsers}$: dummy variable equal to 1 for the users whose ID is smaller than 16,000, as they performed an unusually high number of actions with zero fees;

- $\mathrm{Matchers}$: dummy variable equal to 1 for the legs whose complementary leg was executed by one of the aforementioned 'Anomalous Users';

- Willy: dummy variable for the User ID associated to *Willy*;

- Markus: dummy variable for the *Markus*' User ID not associated to manipulations;

Models (1) to (3) have very low pseudo-R^2; the largest part of the explanatory power of the model is associated with the variables that contain information on the Users IDs. Coefficients β_1, β_2, and β_3 are very small, suggesting that their effect is significant but limited. As expected, instead, the variables $\mathrm{AnomalousDays}$ and $\mathrm{AnomalousUsers}$ are associated with negative and high coefficients. The probability that fees are paid decreases for the users defined as early adopters. Unexpectedly, the variable $\mathrm{Matchers}$ is associated to higher probability that the order might have paid a fee.

Chapter 3

Triangular arbitrage

In this Chapter I mine the leaked history of trades on Mt. Gox to detect the triangular arbitrage activity conducted within the platform. I implement an algorithm to detect the arbitrage actions, then I identify the arbitrageurs and describe their behavior: the availability of user identifiers per trade allows to observe the historical record of each investor. I begin by showing that a considerable difference appears between users when indicators of expertise are taken into account, e.g. whether they conducted arbitrage in a single or in multiple markets. Using this element as a proxy for trade ability, I find that arbitrage actions performed by non expert users are on average non profitable when transaction costs are included (in a mechanism whose dynamics recall the monetary illusion phenomenon), while skilled investors conduct arbitrage at a positive and statistically significant premium. Next, I exploit within-user (across hours and markets) variation and document that expert users make profits on arbitrage by reacting quickly to plausible exogenous variations on the official exchange rates. As a robustness check, I complement the analysis with further indicators of expertise, and I show that results are consistent. A small subset of arbitrageurs is responsible for the vast majority of the arbitrage actions and systematically yields higher profits: I argue that such differences are chiefly due to a better ability of the latter in incorporating information on the exchange rates volatility, which eventually results in a better timing choice at small time scale intervals. The results support empirically the financial literature stating that arbitrageurs are few and sophisticated users.

3.1 Introduction

Arbitrage, the simultaneous purchase and sale of the same asset in two different markets for a risk-free profit, is a key concept in economics and finance. The concept is so important because the *absence* of arbitrage opportunities is a necessary condition for market equilibrium (Harrison and Kreps, 1979). Intuitively, whenever an arbitrage opportunity emerges, some arbitrageur will exploit it until the mechanism of supply and demand has eliminated the price difference. This 'law of one price' makes the no-arbitrage principle a powerful solution concept in financial theory. It is a common foundation of the Capital Asset Pricing Model (Sharpe, 1964; Lintner, 1965; Mossin, 1966), the arbitrage pricing theory (S. A. Ross, 1976), the theory of option pricing (R. C. Merton, 1973; Black and Scholes, 1973), the efficient market hypothesis (Fama, 1970), and many other theories.

In practice, arbitrage is never risk-free. Since purchase and sale are not executed in an atomic[1] transaction across markets, the arbitrageur bears the risk of incomplete execution or concurrent price changes. Moreover, the asset traded in both markets may not be exactly the same, and there may be political risk premia if the markets operate in different jurisdictions (Aliber, 1973). These risks, in addition to other certain transaction costs, impose a lower bound on the price difference needed for profitable arbitrage. The orthodox economic response, in line with the efficient market hypothesis, is to imagine that many small arbitrageurs each take an infinitesimally small portion of the risk (and hence profit). However, Shleifer and Vishny (1997) challenge exactly this view in their landmark work on practical arbitrage in financial markets:

> "[A]rbitrage is conducted by relatively few professional, highly specialized investors who combine their knowledge with resources of outside investors to take large positions." (p. 36)

The authors support this claim by referring to the bounded rationality of many investors, "millions of little traders are typically not the ones who

[1] Incidentally, this may change with decentralized exchanges on programmable cryptocurrency platforms.

have the knowledge and information to engage in arbitrage." (p. 36) While this is plausible and likely cross-checked by expert market participants, the evidence remains anecdotal. Most surprising is the fact that 20 years after these statements were published, I still could not find any academic paper that provides an empirical answer to the question 'Who are the arbitrageurs?'[2], despite the topic being perceived as an issue of compelling relevance in current research (Gromb and Vayanos, 2010). Indeed, in the best case scenario, a comprehensive answer to this question may not only reconcile economic theory with the reality on financial markets, but also refine the assumptions about arbitrageurs in theoretical studies that derive optimal trading strategies in the presence of arbitrageurs (Moallemi, Park, and Van Roy, 2012, and the works cited therein), and ultimately contribute to understanding why persistent limitations to arbitrage emerge.

In this work I seek to provide a partial answer from a very singular market, namely the exchange market between convertible currency and cryptocurrency in the early years of Bitcoin, by mining the leaked dataset of individual and identified trades from Mt. Gox. The choice of market and time is opportunistic. Focusing on the triangular arbitrage activity, I first quantify its magnitude within the Mt. Gox platform, and then I introduce proxy measures for the trading ability of the investors, to show that the expert users conduct more profitable actions, and are more responsive to exogenous shocks on the official exchange rate: a statistically significant relation exists between the investors' expertise and the profitability of arbitrage. my results confirm the anecdotal evidence that arbitrageurs are few and sophisticated users.

The contributions of this work are manifold. First, by investigating the dynamics of the Mt. Gox platform, I enrich the existing literature that focuses on the economic role of cryptocurrency exchanges in the Bitcoin ecosystem (e.g., Moore, Christin, and Szurdi, 2018; Griffin and Shams, 2020). My analyses provide several insights on the Mt. Gox market structure and on its internal trading mechanisms, contributing in this sense

[2]Of course, I simply may have missed the relevant source and would be grateful for pointers from readers.

more broadly to a better understanding of the microstructure of the cryptocurrency markets. Besides this, I provide a detailed estimation of the explicit transaction costs borne by the Mt. Gox users when trading within the platform (while authors previously mainly focused on the analysis of the fees paid for transactions on the Bitcoin network; see e.g. Möser and Böhme, 2015; Dimitri, 2019; Easley, O'Hara, and Basu, 2019). Further, to preprocess the leaked dataset I propose my own deduplication method, that grounds upon (and partially improves) existing methods introduced in the literature (Scaillet, Treccani, and Trevisan, 2017; Feder, Gandal, Hamrick, and Moore, 2018; Gandal, Hamrick, et al., 2018), and whose validity is documented extensively by comparisons with external sources of information.

Second, the availability of user-specific labels enables the design of novel methodologies to investigate arbitrage, based on the analysis of identified sequences of trades. Indeed, previous empirical research focuses mainly on assessing the validity of consolidated theories on arbitrage (e.g., Roll and S. Ross, 1980; Malkiel, 2003), or on detecting violations and anomalies arising in contrast with them (e.g., Lamont and R. H. Thaler, 2003a); however, most of these works are based on aggregate information and do not tackle the topic from a bottom up approach. Other studies do exploit transaction-level data, but lacking the user identifiers, and thus focus primarily on inferring information on the market (C. M. Lee and Ready, 1991) or on the type of trader, rather than on the individuals features (C. M. Lee and Radhakrishna, 2000). The work by Wang et al. (2021), closest to my investigation both in intentions and method, shows clearly that the cryptocurrency ecosystem unveiled unprecedented opportunities in quantitative finance analysis: the authors identify the cyclic arbitrages executed in Decentralized Exchanges (DEXs) within the Ethereum ecosystem. However, the pseudonimity guaranteed by the Ethereum protocol limits the possibility to trace users exactly. In my framework, instead, each leg of all trades is labeled by a user identifier, allowing to trace the sequences of actions executed by the same user and to identify potential links among them. Since in Mt. Gox the bitcoins could be traded within the same exchange against different fiat currencies, my method detects *exactly* the actions of arbitrage as pairs of legs satisfying the textbook properties of arbitrage, that is, two legs (from different trades) executed by the same user, in different currency markets, and within a reasonably small neighborhood of time and volume. It is precisely in this sense that I intend the term 'exactly': the matching legs are identified based on time delay and volume differences, which are known, and conditional on being executed by the same user, another condition which is known. This crucial aspect gives me the unprecedented possibility to identify *completely* the triangular arbitrage activity exploited within the platform, not as an aggregate but at the level of the individual arbitrage actions executed. To the best of my knowledge, this work represents the first attempt to use this approach in the literature; thus, my work provides a contribution also by proposing new algorithms for quantitative finance who exploit micro level information and whose aim is to investigate the individual trading strategies. Noteworthy, this method is also easily extendable to other contexts requiring a pattern identification through user identifiers.

Third, and most importantly, by identifying and describing the characteristics of the individual investors who conduct arbitrage in the presence of risk, my method provides empirical evidence that helps answering broader questions on the nature of arbitrage. The closest papers I could find that investigate the user behavior in financial markets using user level information chiefly focus on the analysis of the individual traders choices, in order to profile investors in terms of risk attitude (e.g., Clark-Murphy and Soutar, 2005; Kourtidis, Šević, and Chatzoglou, 2011; De Bortoli et al., 2019). However, no previous study investigated exhaustively the research questions that motivated my work. This paper thus aims at filling, at least partially, an existing gap between the theoretical description of the arbitrage activity and the practical evidence from real markets. Once the arbitrage activity is identified, I analyse and describe empirically, rather than anecdotally or using aggregate data, the individual trading patterns: I show that only a restricted group of users (N = 440) performs at least one arbitrage action, and an even smaller group of sophisticated users is responsible for the vast majority of trades; the users in such subset are mostly active in multiple currency markets, rather than in a single one, they conduct complex strategies (i.e., metaorders), and consciously control their aggressiveness on the market (by preferring limit to market orders). Most of all, the arbitrage activity of such sophisticated investors is profitable, both on average and after controlling for user fixed effects, though some of their actions yields losses. Crucially, I acknowledge instead that the arbitrage activity attributable to the non sophisticated users is on average non profitable when transaction costs are included, and that it constitutes a small fraction of the total arbitrage activity identified. This raises a relevant conceptual consideration: while a strict definition of arbitrage foresees the execution of such actions only for profit, arbitrageurs accept that they can incur losses, i.e., expected payoffs positive with probability smaller than one, even under fundamentally riskless conditions (Kondor, 2009). However, in this context the non expert users systematically incur losses when the transaction costs are taken into account, and even if some of them are actually executing actions on average positive, thus conducting correctly arbitrage according to the textbook definition, their contribution is negligible in terms of number and volume with respect to the total activity: arbitrage is carried mostly by few and specialized users. I thus conclude that the investors can be classified into in three categories, the non-arbitrageurs, the non-expert arbitrageurs, and the expert ones, the latter two groups being characterized by specific distinctive patterns which I describe in detail.

I stress again the importance of the user identifiers for my analysis. Not only the methodology is based on algorithms that exploit them to capture pairs of legs likely forming arbitrage actions; their role is so essential also as the devised identification strategy is based on the use of user fixed effects, that allow to rule out the possibility that differences in profitability might be influenced by unobservable user-specific ability.

In summary, in this paper I identify exactly the triangular arbitrage activity and the users who conduct it, and I provide the first form of empirical evidence to a relevant issue which up to now was only acknowledged anecdotally: *who are the arbitrageurs?* The answer to this question is relevant in the literature and could help reducing the gap that currently exists between theoretical and practical arbitrage. I point out that my findings may be contingent to this specific market, and that it is not trivial to extend such results to other financial markets, both for the uniqueness of the exchange platform itself and for the specific features that characterize the cryptocurrency market as a whole. Nonetheless, I believe that the general picture I describe reflects well also the features of practical arbitrage in a traditional financial market.

The remainder of the paper is organized as follows. In Section 3.2 I provide some context on the cryptocurrency exchanges, and especially on the Mt. Gox platform. In Section 3.3 I describe the methodology implemented to identify the triangular arbitrage activity and to measure its profitability; I then report descriptive statistics on the users' trading patterns heterogeneity. In Section 3.4 I further explore such differences and provide preliminary evidence of the relationship between the user trade ability, captured by user-specific proxies of expertise, and the profitability of arbitrage. Next, in Section 3.5 I investigate the responsiveness of arbitrageurs to plausible exogenous variations of the official exchange rates, and show that the expert traders react better to sharp variations (in terms of profitability of arbitrage). In Section 3.6 I discuss my findings. Appendices A.1, A.2, A.3 provide additional robustness checks, while Appendix A.4 outlines a description of the arbitrage activity on the major currency market (i.e., the EUR/USD market) taken individually. Appendix A.5 reports supplemental figures and tables.

3.2 Background

3.2.1 Arbitrage in cryptocurrency markets

Arbitrage is a founding and unifying concept in financial economics[3]. It conveys a simple yet powerful message, with vast implications on the theory of asset pricing: assuming that agents are rational and prefer more to less, the no-arbitrage condition states that a portfolio yielding non negative payoffs must have a non negative cost. Excluding the case with zero prices, this statement is equivalent to say that a unique vector of strictly positive state prices, defining unambiguously the price of all assets, indeed exists. When this condition is not met, then there must be a portfolio yielding non negative payoffs without requiring any initial investment - that is, a profitable and riskless investment, which is precisely the definition of arbitrage opportunity (S. Ross, 1976; S. Ross, 1978; Varian, 1987; Dybvig and S. Ross, 1989).

A more recent stream of research reports however evidence that market anomalies arise in the form of persistent mispricings (Lawrence Harris and Gurel, 1986; Froot and Dabora, 1999; Lamont and R. H. Thaler, 2003b), despite the presence of sophisticated investors seeking to exploit such opportunities: arbitrage entails risks and costs and de facto it is often limited. E.g., in De Long et al. (1990) irrational noise traders can operate on an optimistic or pessimistic bias, creating a risk that prevents arbitrageurs from exploiting mispricings and causing price divergences from fundamental values even in the absence of fundamental risk; Shleifer and Vishny (1997) show that when arbitrageurs operate in a principal agency relationship, principals may mistakenly evaluate the arbitrageurs on the base of their performance and refuse to provide the required capital, especially in the extreme circumstances where additional capital would be

[3] In the words of S. Ross (1978), "it is surprising how much of what is central to modern finance is based solely on the arbitrage principles embodied in the basic valuation theorem" (p. 471). Many of the most influential works in finance ground on and contributed to the development and formalization of this fundamental concept. E.g., Modigliani and M. H. Miller, 1958; Black and Scholes, 1973; R. C. Merton, 1973; S. A. Ross, 1976; S. Ross, 1978; Cox and S. A. Ross, 1976; Cox, S. A. Ross, and Rubinstein, 1979; Harrison and Kreps, 1979; Harrison and Pliska, 1981; Kreps, 1981; Delbaen and Schachermayer, 1994.

most needed. R. Merton (1987) points out that markets are not frictionless and investors need to account for additional costs (in capital requirements, entry, strategy implementation, non-istantaneous diffusion of information)[4].

The nascent cryptocurrency market represents a promising area of research in this sense, as its innovative and unconventional market design provides an alternative standpoint to assess the validity of established methods and theory from traditional finance. Bitcoin, the most prominent cryptocurrency in terms of market capitalization[5], is a decentralized system which records transfers between parties denominated in bitcoin (units of cryptocurrency) in a public ledger. By contrast, exchanges are centralized entities in the Bitcoin ecosystem that provide interfaces to conventional payment systems by allowing its users to trade units of cryptocurrency against fiat money (Böhme et al., 2015). Typical exchanges manage and match orders in a private limit order book, and update their customers' account balances in cryptocurrency or fiat money when trades are executed. As a result, exchanges are the place where price formation occurs. Trades on exchanges are kept in a private ledger and have no effect on the public ledger unless users withdraw cryptocurrency from the exchange to a wallet under their own control. Most exchanges publish aggregate information about prices and volume, but concerns about data quality exist since few cryptocurrency exchanges are regulated and audited by the standards of conventional financial markets (Underwood, 2018; Anderson et al., 2018).

I describe here two strategies to exploit arbitrage that are relevant for the cryptocurrency markets and for my context: arbitrage across exchanges and within the same exchange. In the first case, intuitively, an arbitrageur observes a price difference between two exchanges, buys a bitcoin at the cheaper place, then transfers it to the more expensive place, where the bitcoin is sold for a profit. The transfer of bitcoins between exchanges would be observable in the public ledger and could be as-

[4]A survey covering the literature on the limits of arbitrage can be found in Barberis and R. Thaler (2003) and in Gromb and Vayanos (2010).

[5]Around 700 billion $ at the time of writing (according to Coinmarketcap.com `https://bit.ly/3iXhZnj`).

sociated in time with differences in published prices, thereby generating evidence for arbitrage. Indeed, the first approach I devised to detect arbitrage involved the identification of user identities in the Bitcoin network (see, e.g., Meiklejohn et al. (2013)) and the analysis of the on-chain bitcoin flows between exchanges and their users. However, bitcoin transactions are too slow and (at times) too costly for this strategy. Instead, arbitrageurs must maintain a stock of both bitcoins and fiat money in accounts at each exchange in order to react quickly to price differences. The funds can be balanced at a lower frequency and not necessarily correlated with observable price differences. Therefore, while arbitrage opportunities are measurable from published data, there is no way to identify arbitrage transactions or arbitrageurs from the public ledger, and in this regard cryptocurrency exchanges do not offer researchers any advantage over foreign exchange markets. This fact was already mentioned in the first attempt to study arbitrage in Bitcoin (Petrov and Schufla, 2013). Interestingly, this term paper documents that market participants actively explored arbitrage opportunities in spring 2013[6] by pointing to two websites that track suitable price differences (see Figure 37 in Appendix A.5), and one open source software trading bot that exploits arbitrage opportunities. In the simplest case, this strategy of arbitrage requires an investor to open accounts at two exchanges; consequently, to identify arbitrage activity, information from the private ledgers of both the two exchanges involved is needed.

Besides this form of two-point arbitrage, the second relevant strategy is triangular arbitrage: most of the cryptocurrency exchanges offer the possibility to trade bitcoins (or other cryptoassets) against more than one fiat currency. Using bitcoin as a vehicle currency, investors can compare the implied relative price of traditional currencies to the official exchange rate and look for the presence of mispricings. I restrict my analysis to the second strategy, as this form of arbitrage has the advantage that information required to detect the arbitrage activity is entirely contained in the private ledger of a single exchange.

[6]See also, e.g., `https://bit.ly/2GXhSer` and `https://bit.ly/2FwHNJk` for anecdotal evidence.

3.2.2 Relation to Prior Work

Figure 17: Related work in temporal and market context

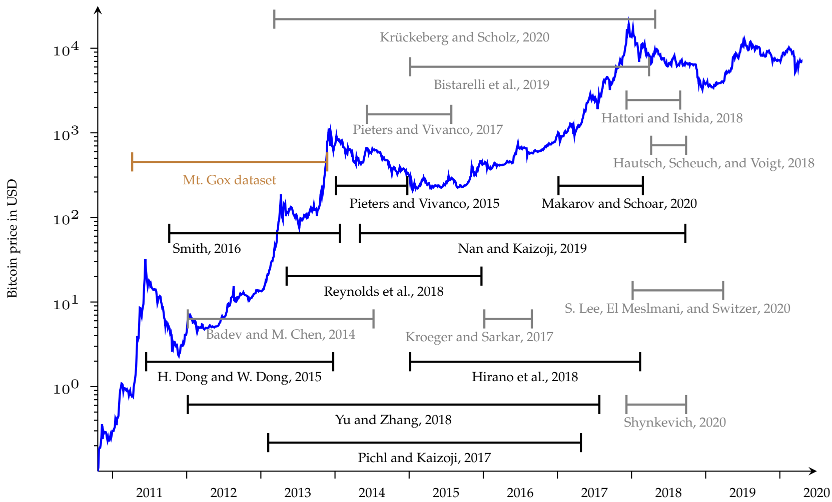

Notes: Bitcoin price in USD on log scale. Black shaded works focus on triangular arbitrage. The gray ones study other forms of arbitrage within the Bitcoin ecosystem.

The body of prior works applying financial econometrics to time series data from cryptocurrency exchanges is vast and not easy to navigate[7]. When interpreting it, it is important to keep in mind that most studies use rather short and often non-overlapping samples. The maturing market for cryptocurrencies has exhibited extraordinary volatility as it transitioned through several epochs. Consequently, the time-series contain multiple structural breaks, which make it hard, if not impossible, to draw conclusions that generalize to the cryptocurrency as a whole. To illustrate this, I depict in Figure 17 the sample periods of the works discussed in the following, along with the bitcoin price in USD on a loga-

[7] e. g., Glaser et al., 2014; Garcia et al., 2014; Brandvold et al., 2015; Yermack, 2015; Cheung, Roca, and Su, 2015; Ciaian, Rajcaniova, and Kancs, 2016; Athey et al., 2016; Bouri, Azzi, and Dyhrberg, 2016; Katsiampa, 2017; Wheatley et al., 2018; Dyhrberg, Foley, and Svec, 2018.

rithmic scale: this accounts best for the order of magnitude differences between epochs. I restrict my review to focused studies of arbitrage opportunities and the exploration of market imperfections. The latter are relevant for my method because they inform about transaction costs, which constrain the exploitability of apparent arbitrage opportunities.

Triangular arbitrage in the Bitcoin ecosystem has been widely investigated by several authors: the early years of Bitcoin trading, which are relevant for my work, are covered by H. Dong and W. Dong (2015) and Smith (2016). The former test the bitcoin market for the presence of triangular arbitrage opportunities between the main cryptocurrency exchanges and the spot currency markets, and exploit price decomposition methods to study to what extent bitcoin behaves as a currency or as a financial asset. They find evidence of persistent price deviations, and conclude that the observed arbitrage stickiness can be explained only if users treat bitcoin as a financial asset rather than as a currency. The latter focuses on Mt. Gox aggregate data, and besides showing that shocks in that market do not affect rates in conventional venues, the author exploits the implied exchange rates in the market to conclude that bitcoin has gold-like (rather than currency-like) properties. He also suggests that there must exist a group of investors who enforce price convergence through their activity, given the degree of efficiency observed in the market. Pichl and Kaizoji (2017) and Reynolds et al. (2018) study triangular arbitrage in a similar epoch, between 2013 and 2017: they both exploit data from Bitcoincharts.com against daily currency spot rates, and find evidence of unexploited arbitrage opportunities. The former investigate triangular arbitrage without considering transaction costs and find unexploited opportunities especially in the Chinese market, while the latter report that in the bitcoin exchanges in exam the persistent mispricings arise only when Bitcoin is used as a vehicle currency, while no deviations from parity arise when considering the rate implied between traditional fiat currencies. Pieters and Vivanco (2015) observe that, from January to December 2014, there are triangular arbitrage opportunities between different exchanges on the USD and EUR markets. As a noteworthy detail, they remark that the exchange rate for the Argentinian pesos (ARS) on

Local Bitcoins, a peer-to-peer exchange, is closer to the ARS black market exchange rate than to the official one. Related to that, Makarov and Schoar (2020) and Yu and Zhang (2018) report evidence of unexploited arbitrage opportunities and suggest that capital controls played an essential role in causing market frictions. Hirano et al. (2018) deploy machine learning techniques to test the Efficient Market Hypothesis in the Bitcoin market, showing that information inefficiencies arise in the form of triangular arbitrage opportunities, especially in minor currency markets. Finally, Nan and Kaizoji (2019) analyse triangular arbitrage on the EUR/USD/BTC currency markets against the EUR/USD FX spot market using FX futures contracts to hedge risk, and show that arbitrage is a competitive strategy.

Beside this, also arbitrage on the same rate but across markets (two-point arbitrage) is well investigated. Empirical studies obtain similar conclusions on the presence of arbitrage opportunities and agree that price deviations, even in different time epochs, emerge and are persistent (Badev and M. Chen, 2014; Pieters and Vivanco, 2017; Kroeger and Sarkar, 2017; Krückeberg and Scholz, 2020); the evidence of mispricings across markets is found also in works proposing theoretical models on arbitrage which are fitted on empirical data (Hautsch, Scheuch, and Voigt, 2018; Bistarelli et al., 2019). Other recent studies focus instead on the nascent futures market for Bitcoin: Hattori and Ishida (2018), Shynkevich (2020), and S. Lee, El Meslmani, and Switzer (2020) obtain partially contrasting findings on the efficiency of such markets (specifically, the first two sources find evidence of efficiency in the markets; the disagreement with the tenor of most other literature can be attributed to the time window and the fact that the futures market operates in a single geographical area). For completeness, I mention that other studies investigate more broadly arbitrage in the cryptocurrency market (e.g., Gandal and Halaburda, 2014; Fischer, Krauss, and Deinert, 2019; Leung and Nguyen, 2019; Crépellière and Zeisberger, 2020).

In summary, based on heterogeneous methods and studying different periods in time with data of different frequency, the literature pretty consistently reports unexploited arbitrage opportunities in cryptocurrency markets. This does not imply that arbitrage does not happen, but might rather indicate that the costs and risks of arbitrageurs are underestimated. Anecdotal evidence from forums, the existence of web-based arbitrage tools, and code repositories for trading bots indicate that arbitrage does happen (Petrov and Schufla, 2013).

3.2.3 The Mt. Gox exchange and the leaked data set

All the above-reviewed studies have in common that they analyze aggregated price (and sometimes volume) time series. My approach differs in that I use individual-level data from the internal ledger of a major exchange, Mt. Gox. The availability of such micro-information is of remarkable relevance: it allows to isolate trends specific to focused groups of traders, as e.g. the arbitrageurs.

Mt. Gox played a prominent role during the early years of Bitcoin: established in 2010, it was the first cryptocurrency exchange and dominated market with around 80-90% of total trading volume until late 2012. It was structured as an order-driven market based on a continuous two sided auction, and formally without any designated specialists. The first competitors entered the market within a short time delay: Bitstamp and BTC-e in mid-2011, BTC China at the end of 2011. Other exchange services entered the market in the following months, and most of them were shut down after a brief period of activity (Moore and Christin, 2013; Ceruleo, 2014).

Since the beginning of Spring 2013, a series of events gradually undermined the Mt. Gox credibility[8], and customers started experiencing delays when withdrawing fiat money[9]. Consequently, the volume traded

[8] 11 March 2013: Mt. Gox temporarily suspends bitcoin deposits after hard fork. `https://bit.ly/2GWPklj`;
11 April 2013: Mt. Gox went down after unexpected increase in the trading activity. `https://bit.ly/3lEuSUW`; 2 May 2013: Coinlab files a lawsuit against Mt. Gox `https://bit.ly/3duT755`; 14 May 2013: the Department of Homeland Security issues a seizure warrant for an account owned by a Mt. Gox's U.S. subsidiary. `https://bit.ly/3k0yhwV`; 5 August 2013: Mt. Gox announces significant losses due to crediting deposits. `https://bit.ly/317TA8n`.

[9] 18 April 2013: users point out withdrawal delays. `https://bit.ly/3lHAyO7`; 4 July 2013: Mt. Gox resumes withdrawals halted on 20 June. `https://bit.ly/3doQxxH`,

in Mt. Gox decreased significantly in the following months, and the bitcoins started to be traded at a large premium in Mt. Gox[10]: in Spring 2013 the competitors of Mt. Gox had already gained a consistent share of the market, pushing Mt. Gox to just under 60% in the summer of 2013. (See Figure 2 in Chapter 1 for market share over time.) The exchange stopped withdraws at once on 7 February 2014, and filed for bankruptcy two weeks later. The former CEO was arrested after criminal charges of fraud and embezzlement in 2015, and found guilty of falsifying data in 2019. Exchange closure is a common phenomenon in the cryptocurrency space, and a source of concern for investors, as witnessed by the survival analysis of 80 exchanges in Moore, Christin, and Szurdi (2018).

The main dataset was leaked to the public in 2014 as a series of CSV files. They contain around 7.5 million trades executed in between 1 April 2011 and 30 November 2013; each trade is composed of two buy and sell legs, reported as separate rows; some information is trade-specific (trade identifier, date of execution, amount of bitcoins exchanged), while other variables are leg-specific (buy or sell type, user identifier, transaction costs paid); further information on the variables is provided in Chapter 2. The vast majority (87.9%) of trades are in USD, followed by EUR (7.7%)[11]. Figure 18 visualizes selected indicators on how Mt. Gox's user base evolved over time, reaching a total of more than 125,000 at the end of 2013. The plots outline intuitively that the peaks of interest towards the cryptoasset (Panel a) and of activity within in the market (Panel b) correspond with periods of exponential growth of the bitcoin price; Panel c sketches the presence of trading patterns across investors who entered in the market in similar epochs: early users tended to operate in the market as sellers, while those that entered from August 2011 to late 2012 are in large part buyers.

I highlight a relevant point for my analysis: whilst the dataset covers a longer period of time, I restrict my sample by excluding the trades exe-

---

`https://bit.ly/3lEvPfY`.

[10]See `https://bit.ly/2FshVy6`.

[11]From September 2011 on, users could trade bitcoins also in exchange for EUR, CAD, GBP, CHF, RUB, AUD, SEK, DKK, HKD, PLN, CNY, SGD, TBH, NZD, JPY (`https://bit.ly/314a5Cg`).

cuted after March 2013, due to the increasing difficulties reported by the users from April 2013 on in managing withdrawal operations and to the other facts documented above. The total users active from April 2011 to March 2013 are approximately 72,000, and the trades are around 5.5 million. Note also that the sample is further restricted, as possibility to trade in currencies other than USD was introduced in September 2011: thus, as Figure 18 shows, the time window considered also coincides with an epoch of constant activity within the exchange platform, and with linear growth rate of the new registered users.

The dataset has been widely analyzed by a number of prior works, which explore research questions that fall apart from arbitrage: e.g., the presence of metaorder executions (Donier and Bonart, 2015), unusual price jumps in the BTC/USD exchange rate (Scaillet, Treccani, and Trevisan, 2017), the effects of distributed denial-of-service (DDoS) attacks on the trading activity (Feder, Gandal, Hamrick, and Moore, 2018), the impact of suspicious activity in the Mt. Gox exchange that likely engaged price manipulation (Gandal, Hamrick, et al., 2018, and W. Chen et al., 2019, the latter through a network science approach), and herding behavior (Haryanto, Subroto, and Ulpah, 2019).

In summary, the dataset is well researched and largely accepted. This, along with my own comparisons to external aggregate information reported in Chapter 2, support its validity and authenticity (and addresses the concerns that may arise to a cautious reader as the source of the dataset is anonymous and there is no way to utterly verify its correctness)[12]. Moreover, according to The Guardian[13], several members of the bitcoin community claimed to have found their own transactions in the dataset. Finally, facts established in the court case against the former Mt. Gox CEO seem to explain patterns in the dataset[14].

I rest on the work of Gandal, Hamrick, et al. (2018), Feder, Gandal,

[12]Comparisons refer to the data published by Bitcoincharts.com. In addition, as in Scaillet, Treccani, and Trevisan (2017), I match my data with an aggregated dataset published by Mt. Gox. The advantage is twofold: I further assess the validity of the leaked dataset, and I collect additional trade-specific information contained only in the latter dataset, i.e., which legs are aggressive (market) orders. Further details are discussed in the Appendix.

[13]`https://bit.ly/2Iu77Rk`

[14]This statement is based on personal communication. I have not read the Japanese files.

Figure 18: Descriptive statistics of Mt. Gox users

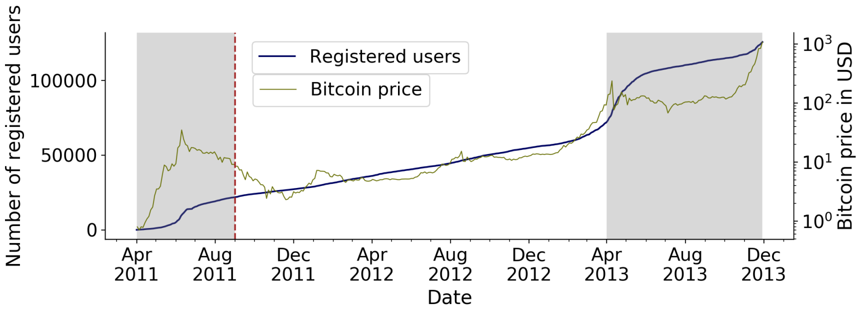

(a)

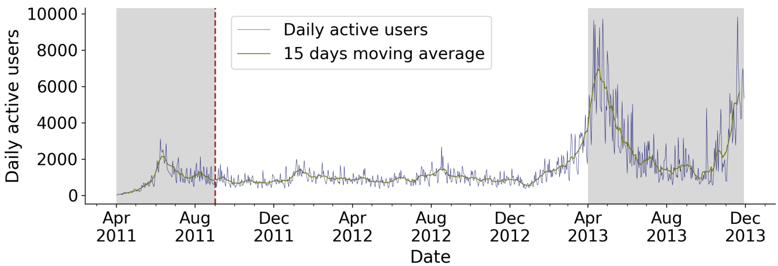

(b)

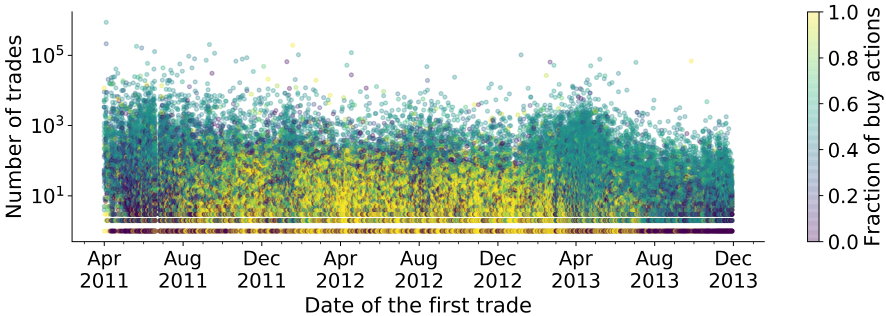

(c)

Notes: Panel (a) shows the growth of registered users in relation to the bitcoin price (the latter on a logarithmic scale). Panel (b) shows the number of daily active users. The brown vertical line indicates the date of introduction of the multi-currency trading; the gray shaded area represents the area excluded by the analysis. Panel (c) shows a scatterplot where each dot represents a user position on the x-axis by the first day of activity on Mt. Gox and on the y-axis by the total number of trades; color indicates the fraction of buy actions.

Hamrick, and Moore (2018), and Scaillet, Treccani, and Trevisan (2017) to pre-process and clean the original leaked dataset. This stage chiefly consists in finding duplicate rows and in identifying (and correcting, when possible) misreported data. The details are described in Chapter 2. It is worth noting that my aggregation technique differs from the above reference in that I aggregate the trades belonging to the same user occurring within the same second. Put it differently, I assume that such actions belong to the same executed order, in compliance with the operating principle of the Mt. Gox filling mechanism[15]. Order speed analyses on other cryptocurrency exchanges reveal that a one-second time scale is suitable to measure order execution delays[16]; traditional financial markets show much shorter latencies (see, e.g., Budish, Cramton, and Shim, 2015; Hasbrouck and Saar, 2013; A. A. Kirilenko and Lamacie, 2015).

Finally, a comment on research ethics and data privacy stands to reason. The internal ledger of Mt. Gox contains data that, in principle, can be linked to natural persons by matching it with other records. Moreover, the users appearing in this dataset had no expectation that their individual trades will become public. I therefore take utmost care that none of my analyses singles out users that have not been singled out in other work (which I always document with a proper citation). Specifically, I map all user identifiers in the dataset to a consecutive sequence of integers, preserving the order but not the numbers. Therefore, user identifiers in my figures should not be directly related to identifiers in the data source. Moreover, I do not possess additional data which would allow linking records to natural persons, nor am I aware of a source where this data could be gathered. Therefore, the harm caused by my study is minimal while there are clear benefits in shedding light into a fundamental question in finance. Readers seeking to replicate the general methods described here are advised to make similar considerations before working with the data.

[15] `https://bit.ly/33YCxaG`.

[16] See `https://bit.ly/3iRJBdu`. Tests show that the time required to add limit orders and execute market orders are comparable, and that execution delays last on average from 10ms to 100/200ms, with skewed tails; a non-negligible amount of trades exists whose latency approaches the second level.

Figure 19: Algorithm to identify the arbitrage actions. Legs executed by user i. Only the green ones form a potential arbitrage action.

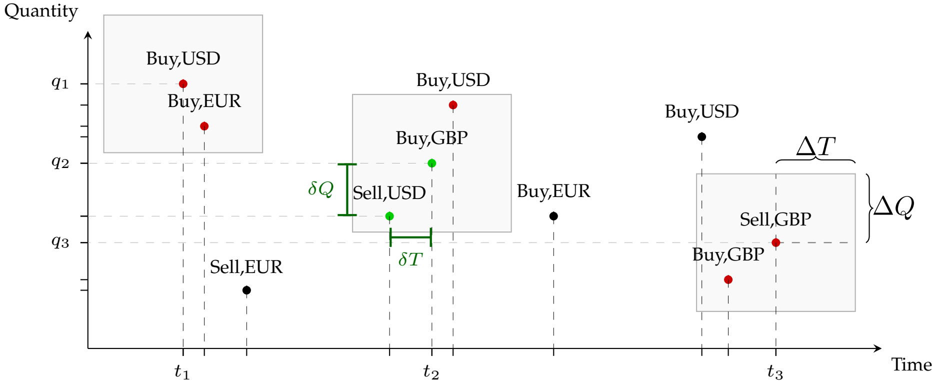

Notes: Each dot represents a leg executed by the illustrative user i, who buys and sells bitcoins (Buy, Sell) against three different currencies (USD, EUR, GBP), as a function of time on the x-axis and volume on the y-axis. The legs are compared only when they are in a small enough neighborhood of time and volume, defined by ΔT and ΔQ. There are three candidate groups of legs: I exclude the legs in the neighborhood of $[t_1, q_1]$, since both are buy legs, and of $[t_3, q_3]$, since the investor trades against the same currency; only the green actions in the neighborhood of $[t_2, q_2]$ form a potential arbitrage action, which is characterized by the values δT and δQ, respectively smaller or equal than ΔT and ΔQ and non-negative by construction. Note that the red dot in the neighborhood of $[t_2, q_2]$ is excluded as it collides with both the two other legs. Note also that a leg can form a single arbitrage action, and if there is more than a matching leg, the closest in time is chosen.

3.3 Identification and description of the arbitrage activity

3.3.1 Arbitrage actions

By definition, triangular arbitrage opportunities in currency markets arise when the exchange rate implied by the ratio of two fiat currencies quoted against a third vehicle currency (in my context, bitcoin) diverges from the official exchange rate. Thus, an investor seeking to exploit such opportunities needs access to at least two currency markets quoted against the same third currency: Mt. Gox users were granted the possibility to trade within the same platform in multiple fiat-to-bitcoin markets, and were entitled to have only one personal account at a time.[17] Bearing in mind that all the legs are labelled by individual identifiers, this setting is ideal to study triangular arbitrage at the micro (individual) level. I thus implement an algorithm to identify arbitrage actions, as described in Figure 19: it is possible to identify a potential arbitrage action in the form of a pair of (buy, sell) legs executed in two separate trades by the same investor using different currencies, and the identification of such arbitrage actions is *exact* - though conditional on the choice of a boundary for the maximum time delay and volume difference between the pair of legs. That is, the triangular arbitrage activity can be observed completely within the private log of one exchange alone.

I focus on the investors that traded bitcoins against more than one fiat currency — only 3,825, out of 71,808; around 1,600,000 legs in the leaked dataset are attributable to them. A subgroup (N = 307) exchanged bitcoins for more than two fiat currencies, being involved in around 800,000 legs. For each leg executed by this group of users I look for potential matches that form arbitrage actions in a neighborhood of time and volume $\pm\Delta T$ and $\pm\Delta Q$: I explore the Mt. Gox log searching for pairs of buy and sell legs that move a nearly equivalent amount of bitcoins, executed (almost) simultaneously by the same user in two separate trades, and exchanged for different fiat currencies — hence in different fiat-to-bitcoin currency markets. I create a data set with all the potential[18] triangular arbitrage actions, irrespective of the currency market they refer to.

In Figure 20 I illustrate how such actions distribute in the space delimited by $\pm\Delta T$ and $\pm\Delta Q$. I present two different cases: in Panel a, where ΔT = 30 seconds and ΔQ = 1%, the number of actions detected is N = 4,464; in Panel b, ΔT = 300s, ΔQ = 10%, and N = 6,629[19]. Some interest-

[17]See `https://bit.ly/2Fu1Rfp` and `https://bit.ly/3j0EkAl`.

[18]Even if these actions respect all the textbook properties of arbitrage, and it is hard to elaborate alternative trading strategies explaining such patterns, I cannot utterly exclude that part of them are false positives. I assume them to be true positives and relax from now on the term *potential*; further details are given in Section 3.6.

[19]Figure 38 in Appendix A.5 shows how the arbitrage actions increase for larger values of ΔT and ΔQ. Table 10 defines formally the parameters ΔT and ΔQ, as well as the action-specific variables δT and δQ.

Figure 20: Distribution of the arbitrage actions by δT and δQ, given the boundaries ΔT and ΔQ.

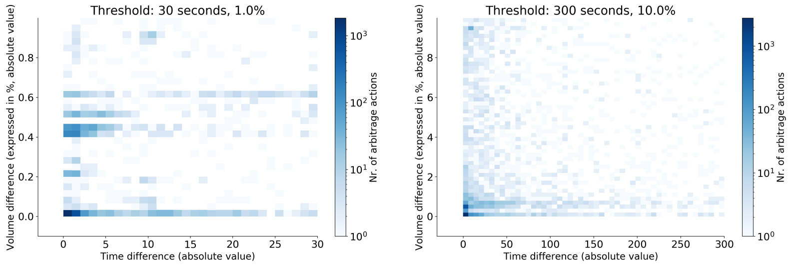

(a) ΔT = 30s, ΔQ = 1%

(b) ΔT = 300s, ΔQ = 10%

Notes: each arbitrage action is characterized by a δT and a δQ representing the distance in volume and time between the two legs composing such action. By construction, they are respectively smaller than ΔT and ΔQ. Both panels report the number of actions in logarithmic scale, as a function of their δT (x-axis) and δQ (y-axis); Panel a focuses on the smaller interval [ΔT = 30s, ΔQ = 1%], Panel b on [ΔT = 300s, ΔQ = 10%],

ing insights come along from this graphical analysis. In particular, as one can see, a large percentage of the trades is observed within small intervals, as the density of actions has a marked peak in the proximity of δT = 0s and δQ = 0%. This does not come unexpectedly, and it matches the textbook definition according to which arbitrage is performed through simultaneous actions involving almost equivalent securities. However, this has also interesting practical implications, as it suggests that likely such precision was achieved via automated trading.

While most of the arbitrage occurred within a tight interval $[\Delta T, \Delta Q]$, I hold a conservative approach and scrutinize in the baseline analysis all the actions occurring within the intervals $[\Delta T = 300s, \Delta Q = 10\%]$, i.e. as depicted in Panel 20b (as a robustness check, in Appendix A.1 I provide the findings obtained on different intervals of $[\Delta T, \Delta Q]$). The vast majority of the 6,629 identified arbitrage actions are performed in the GBP/USD and the EUR/USD markets (respectively 32.2% and 29.1%), followed by EUR/GBP (14.4%), AUD/USD (7.5%), and AUD/GBP (4.7%).

3.3.2 Profitability of the arbitrage actions

Each arbitrage action, which is conducted on a specific fiat-to-fiat currency market, entails some profits (or losses) for the investor, depending on the spread between the exchange rate implied by the same action and the official rate. I then measure the profitability of an arbitrage action as follows:

$$Spread = \frac{ImpER - OffER}{OffER} \cdot 100,$$

where $OffER$ is the hourly official rate[20] and $ImpER$ is the implied one. To compare them, by construction I use the direct quotation with the currency of the buy leg acting as the (fixed) foreign currency. That is,

$$OffER = CUR_B toCUR_S,$$

where CUR_B is the fiat currency used to trade bitcoins on the buy side, and CUR_S on the sell side, and

$$ImpER = \frac{Fiat_S}{BTC_S} \cdot \frac{BTC_B}{Fiat_B}.$$

Noteworthy, the leaked log includes information on the explicit transaction costs incurred by the Mt. Gox users (i.e., the fees associated to each leg of all trades[21]). Although additionally implicit costs may (and are likely to) exist, this feature of the dataset is especially important, as it allows to account for the costs that directly affect the profitability of the arbitrage activity. Thus, in the baseline investigation I adjust the actual profitability by incorporating the leg-specific fees in the prices paid to trade bitcoins, as described formally in Table 10, which provides a recap

[20] I use the hourly open prices of the official exchange rates published on `https://www.histdata.com/`. This dataset lacks information for few minor currencies (CNY, THB, NOK, RUB): as a consequence, I could not calculate the associated profits for 35 arbitrage actions, which are excluded from the analyses where this data is required. E.g., user 5121X in Figure 22d conducted 796 arbitrage actions, but I can calculate the profitability only for 782 of them.

[21] Payable in bitcoins or in fiat currency, and sometimes partly in bitcoin and partly in fiat currency. Users could configure how to pay fees: see `https://bit.ly/34Wyb3h`.

of the main variables described in the last sections. However, as a robustness check, I consider two additional ways to account for the explicit fees (i.e., on a first scenario I exclude them, on another I estimate the fees a user would expect to pay given the official Mt. Gox schedule) which are discussed in Appendix A.2.

The average arbitrage action is worthy of a profit which is 0.42% of the hourly official rate between the two fiat currencies. The average amount of bitcoins traded are equivalent to 52 USD (see Panel A of Table 11 and Table 31 for summary statistics).

3.3.3 A preliminary inspection of the data

Figure 21: Profitability of the arbitrage actions. Users grouped by the number of currency markets exploited for arbitrage.

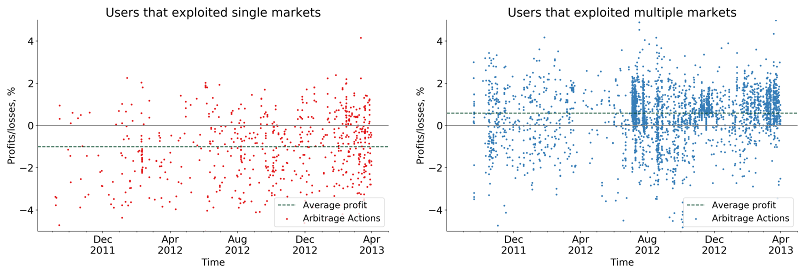

(a) **(b)**

Notes: The plot on the left reports the actions executed by arbitrageurs that exploited a single market. Viceversa, the plot on the right refers to the arbitrageurs who operated on multiple markets. The y-axis reports the profitability of the actions (including fees), depicted as dots, and the x-axis shows their evolution and deployment in time. Note that a negligible number of values may exceed the threshold [-5%,5%] on the y-axis. I do not show them (here and in the following plots) to focus on the area of interest.

The structure of micro-data I collect allows to uncover a number of patterns regarding the behavior and nature of the arbitrageurs. Notweorthy, in disagreement with theory, I note that arbitrageurs are *few* — the set of 6,629 identified arbitrage actions is executed by a total of 440 users (roughly 0.6% of the total users). Furthermore, the arbitrageurs' behavior seems to show a heterogeneous pattern. First, a majority of 395 arbitrageurs explored the presence of opportunities on a single implied fiat-to-fiat currency market — i.e., they exchanged bitcoins for exactly two fiat currencies[22]. Others (N = 45) traded in multiple fiat-to-fiat markets, by exchanging bitcoins for at least three fiat currencies. Remarkable differences appear when comparing the two groups: Figure 21 reports the arbitrage activity of the users who exploited a single market — Panel (a) — and multiple markets — Panel (b). Each dot is an arbitrage action whose x-coordinate is the time of execution and whose y-coordinate is the associated percentage profit/loss. Actions above the gray line are profitable, while actions that lie below determined losses for the users; a dashed line outlines the average profitability. The plots provide graphical evidence that the arbitrage actions executed by users in the latter group are on average more profitable and positive, while the ones in the former are on average negative. Further evidence of such differences is provided by Panels B and C of Table 11, that respectively refer to users who exploited single and multiple markets, and report additional relevant information specific to individual actions, such as the profitability with alternative measurements of the explicit transaction costs, the time delay or the volume difference between the buy and sell sides. They show that the actions executed by users who exploited single markets are on average non profitable, unless fees are excluded, while those conducted by users who exploited multiple markets are always on average profitable. The expected fees instead overestimate the real fees paid for both groups, and the differences between the actual fees paid and the expected fees are larger for the 'Multiple' group. Moreover, the actions in that group are more precise (δT and δQ are on average closer to zero) and, interestingly, are smaller in terms of moved volumes, both considering the amount of bitcoins and of fiat currency. A possible explanation, which I recall below, is that such users exploit more complex strategies

[22] All arbitrage actions involve two fiat currencies traded against bitcoins. Thus, arbitrage actions always refer to a specific fiat-to-fiat currency market. From now on I will imply this concept.

and split orders to reduce the overall market impact.

Second, from Table 11 it emerges also that most of the arbitrage activity is conducted by the users who exploited multiple markets (N = 5,906 against N = 723). Indeed, the three most active users performed 32.8%, 12%, and 10.4% of the total actions, and all of them were active in multiple currency markets. Among those who executed arbitrage on a single currency market, only 11 users performed 10 or more actions, and the most active one performed just 27 actions. Table 12 provides further information on the number of actions executed by the two groups of users.

Third, arbitrageurs that operate on multiple markets are also more acquainted with sophisticated algorithms, such as metaorders. I follow Donier and Bonart (2015) and define as metaorder a group of at least 5 arbitrage actions executed by the same user in the same market (and in the same 'buy/sell direction'), so that the delay between each sequential action never exceeds 60 seconds[23]. As illustrated in Table 13, only 13 arbitrageurs executed metaorders, which are typically composed of less than 10 actions — each delayed of around 20 seconds — and moved an average amount of bitcoins equivalent to few hundreds of dollars. Only 5 users performed more than 5 metaorders, and all of them exploited multiple currency markets and executed more than 100 arbitrage actions. Fourth, arbitrageurs that operate on multiple markets are less likely to behave aggressively. I follow Scaillet, Treccani, and Trevisan (2017) and define the aggressive bids and asks respectively as the buy or sell legs that initiate the market orders. Thus an aggressive arbitrage action is an action with at least one aggressive leg. Table 14 shows that aggressive orders have been used only by users who executed less than 30 actions; on average, they are not profitable.

Noteworthy, clustering the arbitrage actions executed by the same users unfolds interesting insights, and provides further evidence that such differences map into heterogeneous patterns of profitability of ar-

[23]The two parameters are arbitrary. Nevertheless, it is not my main purpose to exactly identify metaorders; rather, I use this measure as an indicator to verify which users exploit these strategies more systematically. As a robustness check, I repeat the procedure increasing the minimum number of actions and the time delay, and the differences are negligible.

Figure 22: Profitability and trading patterns across arbitrageurs

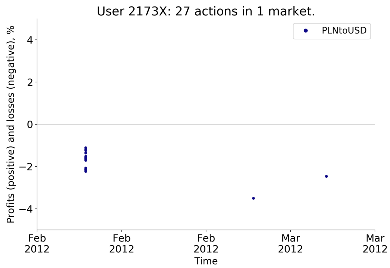

(a)

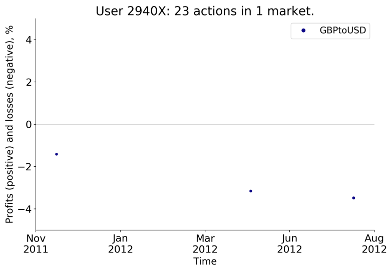

(b)

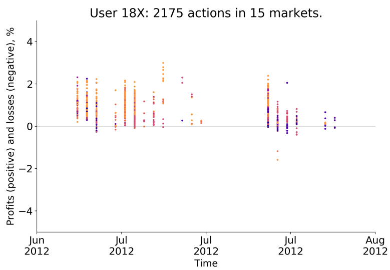

(c)

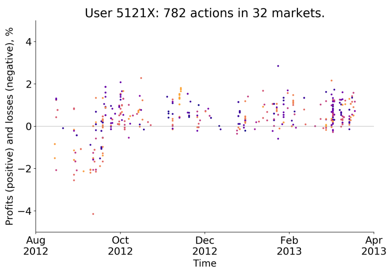

(d)

Notes: the panels report the two most active users in a single currency market (above) and in multiple markets (below). The y-axis indicates the profitability of the actions, depicted as dots, and the x-axis shows their evolution and deployment in time. The different colors correspond to actions conducted in different currency markets. I do not report the legend for the two plots below as the number of markets is too high (15 and 37). I hide the last unit of each user identifier to preserve the anonymity. A negligible number of values may exceed the threshold [-5%,5%] on the y-axis.

bitrage. In Figure 22, for instance, I illustrate graphically the trading pattern of the most active users in a single currency market (Panels a and b) and in multiple markets (Panels c and d)[24]. The dots indicate the profits/losses (y-axis) across time (x-axis) on arbitrage actions. Differences in profits are considerable. While users in (a) and (b) systematically incur in losses when trading (as dots lie below the gray line), the others typically obtain profits by executing far more complex trading patterns. It is also worth remarking important differences between their strategies, e.g. when comparing users 18X and 5121X. Trades performed by the first are concentrated in few weeks (around July 2012); its actions appear as consequential and related and are likely to being part of one or a sequence of metaorders. The trading pattern of the latter is steady and spans across a longer period of time. Both strategies are profitable, non-trivial and likely executed via algorithmic trading.

3.4 Trade ability and profitability of arbitrage

In this paper I hypothesize that arbitrage profitability is a function of the user's trade ability. As laid out in the previous section, my most preferred indicator for trade ability is arbitrage on multiple markets (which I complement with three other variables — number of actions, execution of metaorders and of aggressive orders[25]). Indeed, it is relatively simple to conduct arbitrage exploring opportunities on a single currency market. Evidence suggests that most of the users attempt to do arbitrage in this form (and non-systematically, i.e., in few and dispersed trials, on average non profitable). Few users explore instead more than one market looking for arbitrage opportunities. This activity is in fact far from trivial and requires skills and expertise: users active in multiple markets must set up complex — and likely automated — strategies in order to handle

[24]For completeness, I provide the trading patterns of other traders in 39 Appendix A.5. Furthermore, Figure 26 shows intuitively that a relationship between the variables introduced above and a profitable execution of the arbitrage activity indeed exists.

[25]To further explore the relationship between these variables, I perform a PCA, whose results are reported in Table 15. Table 32 shows instead the Pearson correlation across the main variables of our model.

Figure 23: Days passed between the first arbitrage action of a user and the first one in a new market.

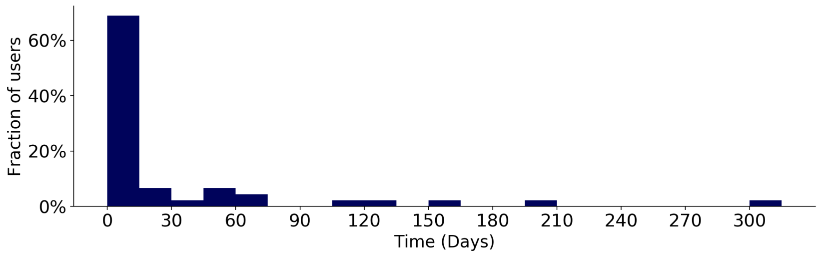

Notes: the graph refers only to the arbitrageurs active on multiple markets. For about 70% of users only 0 to 14 days passed between the first arbitrage action and the first one in another currency market (first bin).

funds in different fiat currencies and to correctly incorporate the increasing amount of disposable information on price variations (the potential number of markets to observe grows non-linearly with the number of currencies used). A further sign of expertise is that operating in multiple markets leads to risk reduction through diversification. Descriptive evidence provided above suggests also that users that make arbitrage through multiple markets obtain on average higher profits. This conclusion is potentially threatened by two facts. First, trade ability may not be fixed but may increase with trading. Second, the correlation between trade ability and profits may be affected by an omitted variables bias. In this section, I take these two aspects into my analysis.

3.4.1 Learning-by-doing and trade ability

The validity of my analysis relies on the fact that arbitrage through trading on multiple markets is a sign of trade ability that a user holds before starting operating on Mt. Gox. Thus my analysis fails to capture the link between expertise and profits if, for example, a user does arbitrage on a single market for long and only after a period of training the user starts arbitraging using other currencies. In Figure 23 I show that such scenario is unlikely to hold in my context. The plot illustrates the distribution of arbitrageurs active in multiple markets across days passed between the first arbitrage action of a user and the first one in a new market. As one can see, the distribution is concentrated in the first bin which gathers arbitrageurs that operate on a new market within 14 days since its first arbitrage action. This bin collects about 70% of them.

In sum, for the vast majority of arbitrageurs the time passed between their first arbitrage action and the first one in a new currency market is short — less than 15 days for 70% of them and a month for about 80%. I therefore conclude that in my context investors sophistication unlikely changed considerably over time. Nonetheless, for robustness I also replicate my analysis excluding users if time delay between the first arbitrage action and the first one in a new market is large, i.e., I remove those that do not fall within the first bin (see Tables 27, 28, 29 in Appendix A.3).

3.4.2 Regression analysis

The difference in profits recorded by users that operate on a single market and users that operate on multiple markets is likely to be biased. For one thing, the latter group may invest a considerable larger amount of money on arbitrage than the former group of arbitrageurs. As the expected profit from trading is larger one may expect that also effort does. For another, profitability may stem from a specific feature of a market or on specific shocks that operate on a single time frame — e.g., external events affecting the volatile Bitcoin ecosystem, sharp price variations and high volatility, and also internal structural changes within Mt. Gox.

A more rigorous way to investigate such difference in profit between the two groups is to estimate the following regression:

$$\mathrm{Spread}_{i,j,p,t} = \beta_0 + \beta_1 \text{Trade Ability}_j + \beta_2\, \text{USD}_{i,j,p,t} + \theta_p + \phi_t + \varepsilon_{i,j,p,t}, \quad (3.1)$$

i indicates arbitrage actions, j users, p the pair of currencies identifying a dyad, and t hours. Residuals, $\varepsilon_{i,j,p,t}$, are clustered at user-level to account for redundant information across actions made by the same user.

The outcome, $\mathrm{Spread}_{i,j,p,t}$, is the distance in price between the implied and the official exchange rate of an arbitrage action i performed by a user j, on the hour t, using a dyad of currencies p. As described in Section 3.3.2, it is reported as a percentage, and by construction the action is profitable when the former is larger than the latter. The explanatory variable of interest in Trade Ability_j is a variable conveying information on the expertise of the user j who conducted the action i. The coefficient of interest is thus β_1, the conditional difference in profits between expert and non expert users (whose profits are captured by the constant, β_0).

Eq. 3.1 also controls for the volume of the trades, expressed in dollars (and divided by 10,000). This variable is preferred to the volume of bitcoins traded as the the latter is subject to high price volatility in time. To construct this variable, prices of the actions not in USD are converted to allow comparisons across currency markets. Most importantly, I include a set of currency pair (dyad) fixed effects, θ_p, that allow to compare arbitrage actions operated using the same couple of currencies. I also introduce hourly time fixed effects, ϕ_t. As I have explained in Section 3.2, Mt. Gox operated at the outset of the bitcoin uptake, was the first exchange platform with a significant relevance, and it was hit by several shocks. Time fixed effects allow to absorb any potential shocks occurred on the market. Besides this, by comparing arbitrage actions filled in the very same hour, I likely capture contingent conditions of the market strictly related to risk, such as liquidity, volatility, depth of the market, that otherwise would be hard to capture given the 'two-leg' (and 'two-currency') structure of the arbitrage actions considered in the analysis.[26]

In Table 16 I present my estimations results where trade ability is proxied by the dummy variable $D(Currencies)$, equal to 0 if the user conducted arbitrage in a single currency market, and 1 if arbitrage is conducted in multiple markets. Overall, these estimations are statistically significant and corroborate the hypothesis that sophisticated arbitrageurs trade on average at a positive premium, relative to less sophisticated users. Namely, column (1) reports the estimate of the correlation between profitability and expertise; in columns (2) and (3) I add separately time and dyad fixed effects; in column (4) I add both fixed effects in

[26] I selected this time scale as a result of a trade-off between granularity and feasibility of the analyses (a smaller scale would be too demanding for a FE-based analysis).

the regression. Some observations are relaxed when including the fixed effects, either because in some hours one single trade was executed, or because a trade is the only one executed in a minor market.

The effect is also economically relevant. Focusing on column (4), I find that the average sophisticated user traded at a premium of 1.292%, relative to the unsophisticated arbitrageurs — a difference which is slightly above a standard deviation in profitability. Finally, it is worth noting that the constant term is systematically negatively estimated across the four specifications. This result is of particular interest if compared to Table 24: once I repeat this analysis without considering the transaction costs, the constant term is non-negative (not significant). My interpretation is that the non sophisticated users do not account correctly such costs, in a mechanism akin to the one originating the monetary illusion phenomenon (Shafir, Diamond, and Tversky, 1997), and thus incur in unprofitable activity. Figure 24 provides an illustrative example for two different time windows: each action is reported twice, once without the transaction costs (and slightly transparent) and once including the fees paid. Red dots represent the actions executed by investors who executed arbitrage in a single market, while the blue ones are executed by expert investors active in many markets. The x-axis is the time of execution of the action, the y-axis is the profit/loss. All these arbitrage actions are affected by the transaction costs, which reduce the yielded profits. However, in the case of the non expert users, not only the actions are in general less profitable, they become even unprofitable once the transaction costs are included, both when they are similar across users (Panel 24b), and when they vary across them (Panel 24a).

Table 17 employs alternative explanatory variables for trade ability described in Section 3.3.3. Column (4) of Table 16 is reported in column (1), for easiness of comparison. Column (2) reports estimations of β_1 when trade ability is proxied by the logarithm of the number of currencies used; in (3) I explore the effect of the logarithmic number of arbitrage actions executed by the user on the profitability of the action. Both estimations are positive and statistically significant. In column (4) trade ability is proxied by the dummy variable $D(Metaorder)$,

Figure 24: The 'monetary illusion' effect.

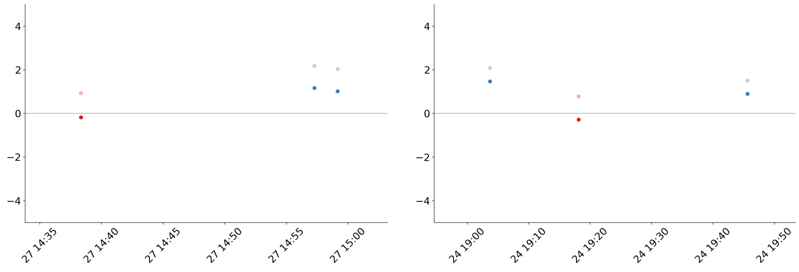

(a) Date: 27 October 2012, h.14

(b) Date: 24 March 2013, h.19

Notes: The plots show the arbitrage activity in two different illustrative time windows. Each action is reported twice, once including and once excluding the transaction costs (the former is transparent, in order to distinguish them). The y-axis reports the profits/losses, and the x-axis the date of execution. In both cases the non expert users (in red) conduct less profitable activity: once the fees are included, their actions become even unprofitable.

which is equal to 1 for *all* the actions conducted by users that executed at least one metaorder. Similarly, in column (5) I use dummy variable $D(Aggressive)$ to indicate whether the user executed or not at least one aggressive action. I stress that both the latter two variables report user-specific and not action-specific information. When trade ability is proxied by $D(Metaorder)$ the sophisticated users are more profitable, but the estimation is less precise; instead, I find that profit reduced significantly when the arbitrage activity is conducted by users who executed at least one aggressive arbitrage action.[27] Finally, column (6) shows the relationship between profits and the scores of the first component obtained with the principal component analysis. Also in this case, the β_1 coefficient is positive and statistically significant. Specifically, I estimate that an arbitrageur with a trade ability score which is a standard deviation above the mean obtained a premium which is around half of the standard deviation in profits (i.e., $\frac{0.224 \times 2.83}{1.26}$).

[27]To better interpret the magnitude of β_1 in Table 31 I report summary statistics for the variables employed.

3.5 Trade ability and responsiveness in arbitrage

The evidence documented thus far suggests that expert users are more likely to make profits on arbitrage relative to non experts. In this section I move on and show that profits stem from a better ability of the former in responding to quick fluctuations which makes arbitrage more (or less) profitable.

A typical profitable situation in financial markets arises when large, unexpected deviations occur in fundamental values. Because of structural frictions, adjustments across markets are not automatic, giving rise to opportunities in conducting arbitrage operations. I exploit this fact in my analysis and reconstruct, from the hourly evolution of the official exchange rate in a market, the percent variation in the exchange rate with respect to the previous hour (see Table 10). The variable $\Delta R_{p,t}$ takes a higher value when the official exchange rate, on a pair of currencies p, observed in the hour t, fluctuates more relative to the previous hour. It is therefore worth remarking that $\Delta R_{p,t}$ varies both across currency markets and time but not within.

The advantage in using this strategy is twofold. First, the exploitation of these temporary opportunities is typically not obvious but requires expertise and/or the execution of automated orders. Hence, when variation in the exchange rate is relatively prominent than usual times it is likely that expert users take advantage of them and make profits. Second, as users who trade in Mt. Gox are small, their actions are unlikely to affect such deviations. It is therefore plausible to assume that users are exchange rate *takers* and deviations in the official exchange rate exogenous.

I employ this variable $\Delta R_{p,t}$ on the right-hand side of my regression and interact it with trade ability to test whether profits obtained by expert arbitrageurs are larger when fluctuations in the exchange rates are larger. This is written as follows:

$$\begin{aligned}\text{Spread}_{i,j,p,t} = &\beta_1(\text{Trade Ability}_j \times \Delta R_{p,t}) + \beta_2 \Delta R_{p,t} + \\ &\beta_3\,\text{USD}_{i,j,p,t} + \alpha_j + \theta_p + \phi_t + \varepsilon_{i,j,p,t}.\end{aligned} \tag{3.2}$$

As one can see, my main variable of interest in Eq. 3.2 is now time variant. This allows to employ a set of user fixed effects, α_j, which permit to absorb any sort of heterogeneity that one may expect across users. This includes education, financial literacy, and other unobservables that are likely to correlate with my measure of trade ability. The inclusion of α_j also implies that my chief variation in the identification of β_1 is the variation across hours (within a user). β_1 can now be interpreted as the difference in profit, between expert and non expert users, following a 1 per cent increase in the rate of change of the official exchange rate. β_2 captures the effect of a 1 per cent increase in the rate of change of the official exchange rate on the arbitrage profits made by non expert users. These effects are additionally identified by including daily time fixed effects, ϕ_t, and currency dyad fixed effects, θ_p, and by controlling for the USD equivalent amount of bitcoin traded, $\text{USD}_{i,j,p,t}$. Standard errors are clustered at user-level as above.

Table 18 reports estimations using different proxies of trade ability. In columns (1-2) I use the variable $D(Currencies)$. I then repeat the analyses by using alternative proxies, namely $Log(Currencies)$ (3-4), the log of the number of actions (5-6), the execution of metaorders (7-8) or aggressive actions (9-10), and finally the scores from the principal component analysis reported in Table 15 (11-12). Overall, I find that an increase in the rate of change of the official exchange rate generates a higher profit for arbitrage made by expert users, even when user fixed effects are included (columns 2, 4, 6, 8, 10, 12). However, β_1 is not statistically significant in columns (2) and (6), perhaps due to the fact that the inclusion of user fixed effects is particularly demanding: as I showed in Table 12, many of the users active in a single market executed just one action. This leads to the exclusion of a large number of observations from this group, making more difficult to obtain stable and statistically significant results. Finally, the result for column (10) — relative to the aggressiveness of the actions — goes against my expectations, but the coefficient is not statistically significant. In support of this analysis I highlight that when I resize the sample as explained in Section A.1 by reducing the boundaries on time and volume, ΔT and ΔQ, I obtain similar results with higher statistical power (see Table 23). Additional robustness checks that further validate my results are reported in Appendices A.1, A.2, and A.3.

In summary, these findings indicate that sophisticated investors are able to take into account and exploit in their favor quick changes in the official exchange rate better than non expert users, and that this ability leads to higher profits. For example, looking at the effect reported in column (12) of Table 18, I estimate that an arbitrageur with a trade ability score which is a standard deviation above the mean obtained a profit which is 1.347% following a 1 per cent increase in the rate of change of the official exchange rate (i.e., 0.476×2.83); note that this premium accounts for more that a standard deviation of the dependent variable. The illustrative examples depicted in Figure 25 provide graphic support and show the arbitrage dynamics at a small time scale: I report the actions executed on the same market, i.e., the EUR/USD one, for four different hour time windows (9 July 2012 h.14, 24 January 2013 h.23, 19 March 2013 h.21, 25 March 2013 h.18 respectively in Panels 25a, 25b, 25c, 25d). Each dot is an arbitrage action, whose y-axis is the profit/loss, and the x-axis is the time of execution. Following the convention introduced above, the blue ones represent the actions executed by the expert users who exploited multiple markets, while the non experts are in red. The gray line represents the absolute value of the minute-level variation of the official rate. The value annotated in proximity of each plot is the specific value of the rate variation in this minute, and it is introduced for the ease of comparison. Note that the value of the rate variation is multiplied by 100. The plots show that users react differently in the presence of exchange rate fluctuations, and are more or less able to turn such changes into their favor. Consider for instance Panel 25c: the expert user conducts arbitrage in a time window with smaller volatility, and its activity is more profitable. A similar pattern appears in Panel 25d, while in Panels 25a and 25b the differences in volatility are smaller. However, also in these cases, non expert users are less able to exploit in a profitable way the variation of the official rate. My interpretation, indeed, is that the expert arbitrageurs are more able with respect to the others to react to price deviations, and thus are also more profitable. In conclusion, thus, I hypothesize that such differences stem from the experts' ability in incorporating information in their strategies (perhaps using APIs, and/or automated trading algorithms); conversely, as low sophisticated arbitrageurs do not incorporate this information correctly, their actions are negatively affected by variations in the official exchange rates. In operative terms, this ability in better capturing information results in a greater expertise at the moment of choosing *when* conducting the arbitrage actions: the timing of execution at the micro scale is a salient element which determines a crucial difference between a profitable and a non profitable action. Thus, ultimately, the differences between the two groups result in a better choice, time-wise, that yields higher profits to the expert users with respect to the non expert ones.

3.5.1 Classification of the arbitrageurs

Thus far I focused on the differences that arise at the level of the arbitrage activity. I first distinguished between non-arbitrageurs and arbitrageurs, and then I investigated the experitse-based differences across users in the latter group. Figure 26 shows intuitively the relationship between the indicators used to proxy expertise and the successful (on average) execution of the arbitrage activity: each dot represents a specific user's average profitability of actions (y-axis) against the number of actions executed (x-axis). The color determines which users exploited one or many currency markets, and the size is proportional to the percentage of arbitrage actions performed through metaorders. I propose here a coarse and illustrative classification, based on the level of expertise, to categorize the arbitrageurs as belonging to the group of skilled or not skilled users, and I briefly describe their overall trading activity (thus including also non-arbitrage trades).

I classify as experts those for which at least three out of the four indicators discussed above are true: executed arbitrage in multiple markets (indicator $D(Currencies) = 1$), conducted a high number of actions (i.e., $Actions \geq 10$), executed metaorders ($D(Metaorder) = 1$), and did not conduct aggressive arbitrage ($D(Aggressive) = 1$). Only 23 out of 440

Figure 25: Illustrative examples of within hours dynamics.

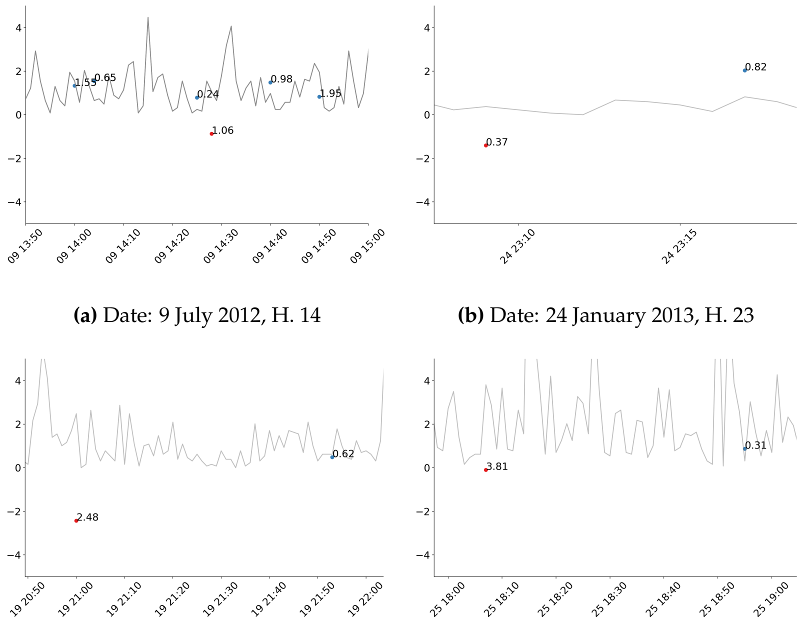

(a) Date: 9 July 2012, H. 14

(b) Date: 24 January 2013, H. 23

(c) Date: 19 March 2013, H. 21

(d) Date: 25 March 2013, H. 18

Notes: each panel reports two plots, vertically stacked, for the same time window. The plots include only actions executed in the EUR/USD market. The actions are dots whose y-axis is the profit/loss, and the x-axis is the date of execution. Red ones are executed by non experts, blue ones by experts. The gray line in the plots below is the percentage variation of the exchange rate with respect to the previous minute (multiplied by 100), and the labels report its value at the minute of execution of each action. The gray line in the plots above is instead the difference between the implied rate and the official rate (at the minute level). The labels corresponding to each action determine the direction of the trade, that is, whether the investor bought bitcoins with EUR and received USD (EUR,USD) or viceversa (USD,EUR).

Figure 26: Users actions: average profits as a function of the number of actions executed.

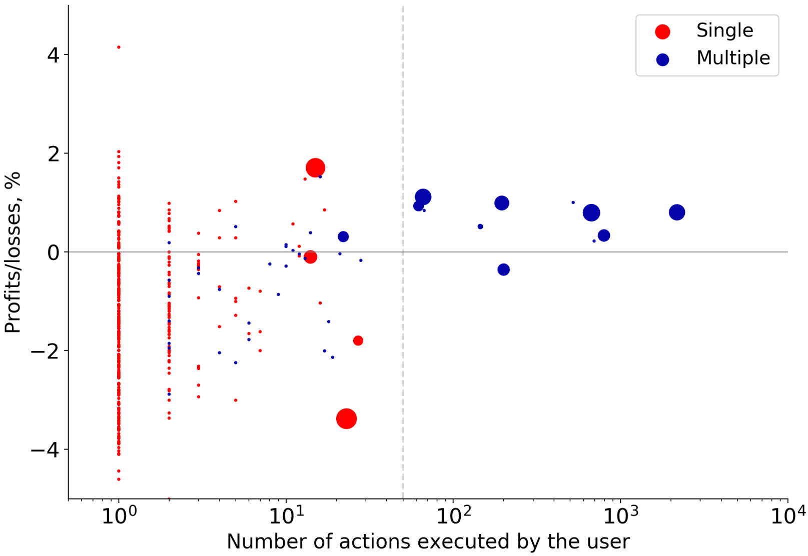

Notes: each dot indicates on the x-axis the number of actions executed by a given user, and on the y-axis the average profits; the blue color is for the users that executed arbitrage actions on multiple markets, while those who exploited a single market are in red; dots whose size is increased correspond to users that executed metaorders, and the size is proportional to the percentage of activity executed through metaorders over the total number of arbitrage actions.

fall within this categorization. Table 19 reports a comparison of the overall trading behavior (which thus includes also the trades not involved in arbitrage activity) of the different categories identified: non-arbitrageurs in columns (a) and (b), arbitrageurs (c), and expert arbitrageurs (f); for completeness I also provide information on the investors who conducted arbitrage in single (d) and multiple (e) markets. Major differences arise across the groups. The characteristics of all the users who traded in Mt. Gox (N = 71,808) are described in column (a); by construction, the expert arbitrageurs in column (e) performed at least 4 trades, thus in column (b) I consider only investors that performed at least 4 actions to balance the comparison. On average, the set of arbitrageurs (c) performed more trades, moved higher volumes of Bitcoins and of fiat currencies, and were active for longer periods of time with respect to the averages on all the users. This difference is evident also in the subgroup (d), and is even more pronounced for subgroups in (e) and (f). Interestingly, the fiat currency moved per trade is comparable across all the groups considered, but smaller both for the arbitrageurs in (e) and (f). This aspect could be partly related to the use of algorithmic trading, and/or the practice of splitting orders to reduce market impact. The User IDs are similar, meaning that on average the users of the different groups entered in the market in comparable periods in time (those in column (f) entered slightly earlier).

3.6 Discussion

The findings in the behavioral finance literature challenge the conventional economic interpretation of theoretical arbitrage that would foresee, in the presence of risk, the intervention of many small traders with homogeneous expectations, not subject to capital constraints, and risk-neutral towards a small enough exposure on the market. Practitioners are well aware that this description is far from reality: arbitrageurs are few, sophisticated, and specialized traders. However, the evidence of this statement still remains anecdotal, and no empirical study exists to describe who are the arbitrageurs. In this work I investigate the Mt. Gox leaked dataset, in order to provide empirical evidence that some of its customers indeed conducted triangular arbitrage activity, as well as an explanation to the observed dynamics. By these means, I try and answer to some open questions on the nature of arbitrage, and on the users who conduct it.

First, I identify the set of potential arbitrage actions, consisting of about 6,600 pairs of buy and sell legs, providing evidence that arbitrage was indeed conducted within the Mt. Gox platform. Coherently with the theory, in the vast majority of cases these pairs are almost simultaneous and involve the same or nearly equivalent security (in terms of volumes of bitcoins moved). The most interesting results concern the characteristics of the investors involved in arbitrage. Out of more than 70,000 users that traded through the Mt. Gox platform in the time window considered, only less than 4,000 traded bitcoins against more than one currency, thus greatly reducing the set of potential arbitrageurs, and just 440 of them are responsible for at least one detected arbitrage action. Major differences across users arise even within this group, if I account for indicators approximating the user expertise. In particular, I focus on the distinction between the arbitrageurs who performed arbitrage only on a single currency market and those who conducted it on multiple currency markets, assuming that it is a good proxy of a trader's expertise: all the members of the former group performed few actions, while the latter, even though the minority, are responsible for the vast majority of actions. They often follow complex trading patterns, including strategies that entail splitting orders into smaller ones to reduce market impact, and operate taking into account the penalizing effect that an excessive aggressiveness would have on their trades. Indeed, the findings are consistent with a scenario in which many uninformed traders move prices far from the fundamental value, and the arbitrageurs absorb the demand shocks by providing liquidity on the market.

I devise a model to quantify the effect of the ability on the profitability of the arbitrage actions: the results show that the actions executed by skilled users are significantly more profitable, while those performed by the non skilled ones are on average non profitable when transaction costs are included. Most of all, the differential effect between non expert and expert users in the presence of large movements of the official exchange rate is positive and statistically significant: the latter are more responsive to exogenous shocks on the official exchange rate. As the essential property of arbitrage is the *advantageous* exploitation of the mispricing of an asset, my findings support the claim that arbitrage is performed by a small number of professional investors.

The results are consistent with the anecdotal evidence that arbitrageurs are few and sophisticated users; however, external validity might be a concern, even in the light of the peculiarities of the Mt. Gox ecosystem. For instance, according to the model described by Shleifer and Vishny (1997), in the conventional financial markets the arbitrageurs exploit their knowledge operating on someone else's funds, thus taking large positions in an agency relationship, where the interplay of capital and risk play a prominent role. In the Mt. Gox market, the average arbitrage action is small (less than a hundred of dollars): I hypothesize that in this context the agency relationship does not take place, and the arbitrageurs are not operating on behalf of someone else. Rather, they are likely investing their own - more limited - funds. This difference could be explained by the fact that I am considering a niche market at its early stages: at the time, the bitcoin market was relatively unknown to the major financial investors. An alternative explanation for the small dimension of the average trade is that expert users implemented also complex strategies involving splitting orders.

The recent findings by Wang et al. (2021) partially mitigate the potential concerns on the external validity. Indeed, they study cyclic arbitrage across decentralized exchanges (DEXs) in the second most prominent cryptocurrency after Bitcoin, Ethereum. The market shows relevant differences, as it is based on the use of smart contracts[28], and because the investigation was conducted on a very different time epoch: they observe almost 300,000 arbitrage cycles between May 2020 and January 2021, a sign that the market has become much more mature with time. Most interestingly for my purposes, they notice relevant differences between two groups of users, those who exploit private smart contracts and those who conduct arbitrage using public protocols, the former experimenting a much higher success rate (62% against 28.4%). Creating private private smart contracts is complex and requires deep knowledge of the Ethereum ecosystem, in order not to incur large losses: this finding, obtained for a different cryptocurrency and in a very different epoch, is strictly related to my study and consistent with it, thus further supporting the validity of my investigation.

Some other limitations to this work stand out: a valid objection is that all these actions could be false positives. It is impossible to prove undoubtedly that the actions are arbitrage, unless a direct proof is provided by the Mt. Gox users (I thus encourage the Mt. Gox arbitrageurs among the readers to contact me and share their comments). However, some major elements emerge in support of my hypothesis. First, it is hard to find alternative explanations to the existence of several pairs of trades executed according to patterns so specific such as those I described throughout the paper. Second, I explicitly consider several variants and controls in my regressions, in order to rule out alternative hypotheses. Third, the fact that many actions - predominantly profitable - are ascribable to few large users is in my opinion another meaningful factor.

Another limitation is that I implicitly assume that I am providing an 'upper bound' to the total triangular arbitrage activity: assuming that the identified actions are true positives, then my algorithm ideally detects all

[28]See XXXX for an introduction on the concept of smart contract.

triangular arbitrage activity. However, this is likely not true: it is possible that some arbitrage actions have a more complex structure than just being composed of two buy/sell legs (e.g., they might involve many cycles and/or several currencies), and that the aggregation method described in Section 3.2 to emulate the Mt. Gox internal matching mechanism is not thoroughly accurate. Thus, even though I believe my procedure is a good approximation, I acknowledge that I might be estimating imperfectly the total number of arbitrage actions. Furthermore, I select the arbitrage actions in a reasonably small neighborhood of time and volume, but I do not ground my choice on a theoretical support; rather, I base it on empirical evidence. For this reason, I adopt a conservative scenario with larger boundaries, and I report additional estimations both on a less conservative scenario and on a case with even larger thresholds in Appendix A.1. I do not investigate the effect of the arbitrage activity in enforcing the 'law of one price': an interesting development of the current work would be to inspect whether, and to what extent, the arbitrage activity had an effect on the market efficiency.

Finally, the analysis focuses on a single exchange platform: on the one side it is an advantage, as all information required to detect arbitrage activity is entirely contained in the private ledger. On the other side, triangular arbitrage only aligns prices in one market, whereas an essential function of arbitrage is its function of 'information carrier' between markets. So, while the evidence might be stronger for triangular arbitrage within the same market, this result may miss generality.

3.7 Tables

Table 10: Definitions.

Variable	Description/formalization
Arbitrage Action	Action composed of two legs executed by the same user in different trades using different currencies. The time delay and volume difference cannot exceed a threshold $[\Delta T, \Delta Q]$. Each arbitrage action is obtained by merging a buy and sell leg. $$ArbitrageAction = (Leg_{Buy}, Leg_{Sell}) \quad (3.3)$$
ΔT	Maximum time delay allowed (e.g., 300 seconds in the baseline analysis)
ΔQ	Maximum volume difference allowed (e.g., 10% in the baseline analysis)
δT	Time delay between Leg_B and Leg_S, expressed in seconds. By definition smaller or equal to ΔT: $$\delta T = \lvert T_B - T_S \rvert \leq \Delta T \quad (3.4)$$
δQ	Volume difference (Bitcoins traded) between Leg_B and Leg_S, expressed as a percentage. By definition smaller or equal to ΔQ: $$\delta Q = \frac{\lvert Vol_B - Vol_S \rvert}{(Vol_B + Vol_S)/2} \cdot 100 \leq \Delta Q \quad (3.5)$$
Official exchange rate	By convention, each arbitrage action is compared to the official exchange rate in the following way: $$OffER = CUR_B toCUR_S \quad (3.6)$$ that is, if the buy leg of an arbitrage action is performed in EUR and the Sell one is in USD, then I consider the official exchange rate EURtoUSD. If the Buy side is in USD, and the Sell one in EUR, then it is compared to the e.r. USDtoEUR.

R — Hourly variation of the official exchange rate expressed as a percentage:

$$R = \frac{|OffER_{t_1} - OffER_{t_0}|}{OffER_{t_0}} \cdot 100 \tag{3.7}$$

Dyad — Pair of currencies that defines the fiat-to-fiat currency market to which the arbitrage action belongs. E.g., the dyad (EUR,USD) refers to actions whose currencies are CUR_B = EUR and CUR_S = USD or viceversa (as they 'refer' to the same currency market).

Implied exchange rate — The implied exchange rate is calculated by comparing the price of bitcoins in the two legs. The latter row includes fees.

$$ImpER = \begin{cases} \dfrac{Fiat_S}{BTC_S} \cdot \dfrac{BTC_B}{Fiat_B} & \text{without fees} \\ \dfrac{Fiat_S - Fee_{f,S}}{BTC_S + Fee_{b,S}} \cdot \dfrac{BTC_B + Fee_{f,B}}{Fiat_B - Fee_{b,B}} & \text{with fees} \end{cases} \tag{3.8}$$

The pedices B and S refer to the buy and sell side; f and b indicate respectively if the term Fee is denominated in fiat or in bitcoins.

Profit (Spread) — Spread between the implied and the official rate divided by the official rate, expressed as a percentage. By construction, profits arise when ImpER > OffER.

$$Spread = \frac{ImpER - OffER}{OffER} \cdot 100 \tag{3.9}$$

Aggressive — Arbitrage action composed by at least one aggressive leg (that is, a leg that initiated a market order).

Equiv. $ — Value of a trade expressed in dollars. I use this variable to indicate the value of a trade since the bitcoin value is not stable in time.

Metaorder	Metaorders are identified as sequences of at least 5 arbitrage actions executed by the same user, in the same market, and such that the time passed between each action is less than one minute. *Note*: I partly follow the methodology described in Donier and Bonart (2015), with some differences: the authors consider a larger time delta (one hour) bewteen each action, and contrary to them I use an arbitrary parameter (N=5) to define the minimum length of a metaorder. While I do not provide the results here, I varied the two thresholds and noticed that the differences are negligible for my purposes, and this setting reduces the false positives classification.

Table 11: Descriptive statistics of the arbitrage actions

	Mean	St.D.	Min	25%	50%	75%	Max
Panel A: all arbitrage actions (N = 6629)							
Profits, fees, %	0.42	1.26	-11.35	0.075	0.621	1.096	18.16
P., exp. fees, %	0.28	1.22	-7.46	-0.191	0.375	0.982	18.24
P., no fees, %	1.05	1.21	-6.40	0.490	1.110	1.696	19.60
Bitcoins	4.12	12.56	0.00	0.039	0.807	3.261	334.14
'Equiv. $'	52.54	169.63	0.00	0.359	7.400	41.424	4666.66
δT (s)	29.04	59.09	0	0	1	24	300
δQ (%)	1.30	2.46	0.00	0.000	0.215	0.863	9.99
Panel B: actions of users who exploited single markets (N = 723)							
Profits, fees, %	-1.00	1.96	-11.35	-2.191	-0.933	0.118	18.16
P., exp. fees, %p	-0.95	1.91	-7.46	-2.161	-0.891	0.172	18.24
P., no fees, %	0.11	1.91	-6.40	-1.135	0.134	1.275	19.60
Bitcoins	7.89	21.16	0.00	0.253	2.000	7.472	288.35
'Equiv. $'	118.06	340.25	0.00	4.014	27.299	95.708	4666.66
δT (s)	59.95	68.21	0	13	34	86	297
δQ (%)	1.04	1.77	0.00	0.461	0.602	0.602	9.82
Panel C: actions of users who exploited multiple markets (N = 5906)							
Profits, fees, %	0.59	1.02	-7.40	0.208	0.688	1.128	10.13
P., exp. fees, %p	0.42	1.02	-7.34	-0.009	0.448	1.019	10.15
P., no fees, %	1.16	1.04	-6.28	0.577	1.178	1.719	10.79
Bitcoins	3.66	10.97	0.00	0.030	0.606	2.995	334.14
'Equiv. $'	44.52	132.48	0.00	0.318	5.767	35.087	3862.71
δT (s)	25.26	56.74	0	0	1	16	300
δQ (%)	1.34	2.53	0.00	0.000	0.000	0.928	9.99

Notes: actions identified at ΔT = 300s and ΔQ = 10%. Panel A describes the main features of all the arbitrage actions, while Panel B reports the statistics for the subset of actions (N = 723) executed by users that performed arbitrage in a single currency market. Panel C refers to those executed by investors active in multiple markets (N = 5,906).

Table 12: Statistics on the number of actions executed by the arbitrageurs.

	Mean	Std	Min	25%	50%	75%	90%	95%	Max
Group *Single* (N = 395)	1	2	1	1	1	2	2	5	27
Group *Multiple* (N = 45)	131	366	2	4	11	28	392	690	2175

Notes: I split the users in two groups, that is, those who performed arbitrage on a *Single* and on *Multiple* markets. The statistics describe the mean, standard deviation, minimum, maximum, and percentiles of the number of actions performed by the two subgroups of users. Note that, by construction, the users in the group *Multiple* performed at least two arbitrage actions; thus, they are involved in at least four trades. Similarly, users in the group *Single* conducted at least two trades.

Table 13: Arbitrage actions executed via metaorders, descriptive statistics.

	Percentage	Number of metaorders	Avg. length	Avg. time delay	Avg. Bitcoins	Avg. Equiv. dollars
18X	54.07	91	12.92	13.33	52.54	369.38
1245X	80.00	2	6.00	23.83	7.43	97.73
1964X	44.10	11	7.82	26.73	35.81	178.17
2173X	18.52	1	5.00	14.00	40.00	234.55
2286X	35.71	1	5.00	26.25	5.00	297.55
2717X	3.45	1	5.00	47.00	0.59	6.45
2940X	91.30	1	21.00	17.75	2.36	25.80
3174X	63.28	40	10.60	29.22	30.81	346.89
4156X	29.00	7	8.29	28.46	9.97	70.54
4325X	22.73	1	5.00	29.50	55.00	1118.86
4901X	56.06	1	37.00	11.36	16.55	162.88
5121X	29.40	26	9.00	15.51	1.32	35.86
6688X	20.97	2	6.50	20.07	7.36	242.74

Notes: for each user (rows), I identify the sequences of actions with the characteristics of metaorders. Only the 13 users reported here performed metaorders. *Percentage* indicates the number of actions that are part of metaorders over the total number of arbitrage actions executed by the user; the second column represents the number of metaorders identified. The other columns describe average values on the metaorders executed by each user and respectively report the average number of actions that compose a metaorder, the average time delay between the actions in the same metaorder, the mean volume of a metaorder expressed in dollars and in bitcoins. User identifiers are anonymized.

Table 14: Descriptive statistics of the aggressive arbitrage actions (N = 313).

	Mean	St.d.	Min	25%	50%	75%	Max
Arbitrage actions (N)	6.572	8.223	1.000	1.000	2.00	11.0	28.000
Spread (%)	-1.106	1.434	-5.354	-2.13	-0.975	-0.058	2.243
$Currencies_d$ (dummy)	0.278	0.449	0.000	0.000	0.00	1.0	1.000

Notes: out of N = 6,629 arbitrage actions, just N = 313 are aggressive actions, that is, arbitrage actions in which at least one of the two legs of the arbitrage action is an aggressive order. They are executed by users who performed few arbitrage actions (1^{st} row: 6.57 on average, and maximum 28); on average they are not profitable (2^{nd} row), and they are executed primarily by users active only on single markets (3^{rd} row).

Table 15: Principal Component analysis.

	$D(Currencies)$	$Log(Actions)$	$D(Metaorder)$	$D(Aggressive)$	Expl. variance, %
PC1	0.54	0.63	0.53	-0.15	53.38

Notes: the sample is the set of users (N = 440). I consider the four main indicators that I exploit in Section 3.5: the dummy variable that classifies users who exploited single or multiple markets, $D(Currencies)$; the logarithm of the actions executed, $Log(Actions)$; whether the action is part of a metaorder ($D(Metaorder)$) or aggressive ($D(Aggressive)$). Values are standardized by constructing their z-score. The table shows the loadings of all the variables (columns) for the first principal components (row). The last column reports the explained variance of the component.

Table 16: Relationship between trade ability and profits

Dep. var.:	Spread (with fees)			
	(1)	(2)	(3)	(4)
D(Currencies)	1.6180***	1.5791***	1.2421***	1.2917***
	(0.1900)	(0.1943)	(0.1584)	(0.1659)
Equiv. $	3.4652**	2.6151*	0.6556	0.1912
	(1.7532)	(1.5862)	(1.2465)	(1.1280)
Constant	-1.0420***	-0.9985***	-0.6141***	-0.6506***
	(0.1593)	(0.1775)	(0.1500)	(0.1604)
Time FE	N	N	Y	Y
Dyad FE	N	Y	N	Y
N	6594	6582	5307	5284
R-squared	0.16	0.20	0.68	0.69

Notes: the Table reports OLS estimates of the relationship between the dependent variable *Spread*, that captures the profitability of an arbitrage action, and the variable D(Currencies), which is a proxy of the user trade ability, equal to 1 if the user conducted arbitrage in multiple markets, and 0 otherwise. I consider four different specifications of the model: (1) without including fixed effects, (2) with dyad fixed effects, (3) with time fixed effects, (4) with both. All columns include an additional control for the amount of volume traded, expressed in USD (and divided by 10,000). I report only the overall R^2. Errors are clustered at the user-level to account for intra-class correlation. Standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 17: Relationship between trade ability and profits, alternative proxies

Dep. var.:	Spread (with fees)					
	(1)	(2)	(3)	(4)	(5)	(6)
D(Currencies)	1.2917***					
	(0.1659)					
Log(Currencies)		0.9326**				
		(0.4439)				
Log(Actions)			0.3165***			
			(0.0627)			
D(Metaorder)				0.2877		
				(0.1914)		
D(Aggressive)					-1.5280***	
					(0.1796)	
PC1						0.2242***
						(0.0466)
Equiv. $	0.1912	0.1525	1.2060	0.0733	0.3375	0.8175
	(1.1280)	(1.2621)	(1.3191)	(1.3384)	(1.1244)	(1.3500)
Constant	-0.6506***	-1.0453	-1.4603***	0.3359**	0.6113***	-1.0717***
	(0.1604)	(0.7807)	(0.4143)	(0.1614)	(0.0424)	(0.3593)
Time FE	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y
N	5284	5284	5284	5284	5284	5284
R-squared	0.69	0.69	0.72	0.67	0.69	0.70

Notes: the Table reports OLS estimates of the relationship between the dependent variable *Spread* and alternative proxies of the user trade ability: (1) D(Currencies) provides a baseline reference by repeating column (4) of Table 16; (2) Log(Currencies) is the logarithm of the number of currency markets exploited by the user; (3) Log(Actions) is the logarithm of the number of arbitrage actions executed by the user; (4) and (5), D(Metaorder) and D(Aggressive), are respectively dummy variables that indicate whether the user conducted metaorders or aggressive actions. (6) PC1 is the score of each variable obtained by performing a PC analysis as explained in Table 15. All columns include time and dyad fixed effects, as well as an additional control for the amount of volume traded, expressed in USD (and divided by 10,000). I report only the overall R^2. Errors are clustered at the user-level to account for intra-class correlation. Standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 18: Dependent variable: Spread, with fees. Interaction between expertise and daily rate variation. Coefficients fitted with OLS.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Int. D(Currencies)	6.905***	1.442										
	(1.139)	(5.314)										
Int. Log(Currencies)			3.213**	1.703**								
			(1.528)	(0.723)								
Int. Log(Actions)					1.427***	0.150						
					(0.202)	(0.409)						
Int. D(Metaorder)							3.693**	2.228***				
							(1.600)	(0.821)				
Int. D(Aggressive)									-6.786***	3.593		
									(1.595)	(3.461)		
Int. PC1											1.048***	0.476**
											(0.159)	(0.211)
DeltaER	-5.441***	-0.742	-4.744**	-2.426	-7.060***	-0.258	-2.341*	-1.045	0.168	0.617	-5.791***	-2.810*
	(1.196)	(5.279)	(2.125)	(1.590)	(1.320)	(2.423)	(1.241)	(0.854)	(0.526)	(0.956)	(1.187)	(1.546)
Equiv. $	1.150	-0.732	0.732	-0.861	2.128	-0.751	0.766	-0.804	0.824	-0.741	2.225	-0.812
	(1.700)	(0.833)	(1.762)	(0.878)	(1.815)	(0.827)	(1.647)	(0.871)	(1.896)	(0.838)	(1.800)	(0.866)
User FE	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Time FE	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
Dyad FE	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y
N	6594	5142	6594	5142	6594	5142	6594	5142	6594	5142	6594	5142
R-squared	0.05	0.75	0.02	0.75	0.07	0.75	0.02	0.75	0.02	0.75	0.06	0.75

Standard errors in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 19: Comparison of identified arbitrageurs to non-arbitrageurs

		All users	Users with ≥ 4 trades	Arbitrageurs			
				All	Single	Multiple	Classified
	N	$71,808$	$47,201$	440	395	45	23
		(a)	(b)	(c)	(d)	(e)	(f)
Total trades	mean	77.79	117.41	2,077.13	600.88	14,707.27	28,173.13
	50%	7	15	83	66	842	2,451
Arbitrage actions	mean	0.09	0.14	15.33	1.85	130.73	251.65
	50%	0	0	1	1	11	27
Currencies	mean	1.06	1.09	2.50	2.18	5.22	4.22
	50%	1	1	2	2	4	3
USD trades	mean	70.63	106.71	1,639.35	468.48	11,656.73	22,643.09
	50%	4	11	39.5	31	420	974
EUR trades	mean	3.51	5.28	216.04	101.55	1,195.58	2,134.96
	50%	0	0	6	3	57	420
Buy trades	mean	38.37	57.73	961.54	258.11	6,979.78	13,340.83
	50%	4	8	32	26	345	1,294
Sell trades	mean	39.42	59.68	1,115.59	342.77	7,727.49	14,832.30
	50%	1	4	45	38	531	1261
Bitcoins bought	mean	532.96	803.79	6,332.91	2,339.71	40,496.94	77,816.93
	50%	18.2	46.2	146.7	123.7	904.3	3917.4
Bitcoins sold	mean	532.96	807.55	7,046.89	3,062.71	41,133.74	78,812.37
	50%	2.0	13.6	186.7	164.4	878.7	4,157.9
'$' sent	mean	6,000.03	9,028.33	84,535.25	30,326.24	548,323.41	1,044,742.87
	50%	237.7	554.6	1,943.0	1,479.7	15,144.1	48,984.1
'$' received	mean	6,000.08	9,075.14	91,647.19	38,076.37	549,975.38	1,046,048.90
	50%	31.3	200.4	2614.5	2010.8	15147.7	33674.9
'$' per trade	mean	158.94	149.97	131.75	134.09	111.75	91.57
	50%	67.1	67.7	64.5	67.5	52.1	54.7
Active days	mean	10.23	14.86	55.17	48.20	114.80	190.65
	50%	3	6	23	20	55	171
Active hours	mean	20.67	30.69	237.14	145.38	1,022.20	1,790.96
	50%	4	8	38	34	200	598
Trades per day	mean	3.16	4.09	10.75	6.43	47.64	79.97
	50%	2	2.5	3.9	3.4	12.1	20.3
Trades per hmy	mean	2.17	2.62	3.18	2.80	6.43	7.97
	50%	1.5	1.9	2.3	2.1	3.9	4.8
User ID	mean	35,903.50	32,745.60	33,816.86	33,726.13	34,593.13	31,252.04
	50%	3590X	3175X	3133X0	3117X	3899X	2763X

Notes: (I) on the meaning of '$'. To make the results comparable, I converted in USD the value of the trades denominated in different fiat currencies. (II) The user ID is a sequential value that does not correspond to the one originally reported in the leaked dataset; the unit is hidden to preserve anonymity.

Chapter 4

Two-point arbitrage in the Bitcoin ecosystem

In this Chapter I identify the investors who performed two-point arbitrage across exchange platforms in the Bitcoin ecosystem. I investigate a time period ranging from April 2011 to March 2013, and focus on the activity in USD between Mt. Gox and two other cryptocurrency exchanges, BTC-e and Bitstamp. This choice is subordinated to the availability of the dataset containing the history of trades executed within the Mt. Gox platform: each leg of all trades is labeled by a user-specific identifier, allowing to reconstruct the sequence of actions executed by each investor. I match these sequences to 'ideal' sequences of arbitrage trades, considering the price differences between exchanges and a user-specific estimate of transaction costs. The latter involves the fee model described in Chapter 2 inspired by the posted fee schedules and fitted to empirical data. The subset of investors whose actual series matches the ideal series best are potential arbitrageurs: out of around 72,000 users, I identify 1,441 potential arbitrageurs with all two counterpart exchanges. I cross-compare their activity within the Mt. Gox platform to the history of trades in the anonymized logs of the counterpart exchanges, and construct a second metric that indicates if matches (equivalent and simultaneous trades) are found. I post-filter the first estimation and further reduce the set of detected users to 10 arbitrageurs with Bitstamp and 45 with BTC-e, for a total of 49. A comparison of aggregate statistics between arbitrageurs and non-arbitrageurs is given and discussed.

4.1 Introduction

Bitcoin is a communication protocol that facilitates the execution of value-based transactions by preserving the user anonymity, and concurrently solving the double-spending problem for a purely digital currency, bitcoin, without resorting to a third party (Nakamoto, 2008a). It provides pseudonymous users (Androulaki et al., 2013) with a virtual infrastructure to broadcast transactions, based on a decentralized peer-to-peer communication network (Ron and Shamir, 2013), and relies on the activity of miners, a subgroup of users who are rewarded in bitcoins - conditional to performing a computational intensive task, the Proof of Work (PoW) - to secure the system and ensure the validity of the proposed transactions (Kroll, Davey, and Felten, 2013; Eyal and Sirer, 2014). This mechanism eases the achievement of consensus across untrusted participants over the state of the system - that is, the agreement on who owns and has the authority to spend an amount of the valuable digital currency, bitcoin - in a decentralized fashion (Barber et al., 2012; Bonneau et al., 2015; Garay, Kiayias, and Leonardos, 2015). Precisely this feature led the academic community, along with practitioners, to envisage in it a potentially disruptive technology (Catalini and Gans, 2016).

Whilst its design principle is non-trivial, requires deep technological skills to avoid operational errors and exposes users to non-conventional risks, Bitcoin started attracting soon the attention of investors not necessarily exposed to such technological knowledge, as an investment asset rather than as an alternative payment system (Glaser et al., 2014; Yermack, 2015). Exchange platforms, which are centralized intermediary services that facilitate the trading of bitcoins against conventional payment systems (Böhme et al., 2015), gained consideration and attention precisely for this reason, as they reduce the barriers to entry into the Bitcoin ecosystem. They rely upon trading mechanisms adopted from traditional finance, and behave similarly to traditional banks (Anderson et al., 2018). They act de facto as intermediaries between the Bitcoin network of users and the traditional finance system.

Such exchanges have a relevant role in the Bitcoin ecosystem in terms of volumes moved (Lischke and Fabian, 2016), and ultimately are the place where the price formation mechanism takes place. Noteworthy, the internal dynamics of different bitcoin exchanges are not directly related. An essential question is how such exchange-specific information is broadcast across different platforms, and to what extent prices in other venues embody it: the founding financial concept of market efficiency foresees that in an efficient market the asset price incorporates all relevant information (Fama, 1970). The concept of market efficiency is strictly connected to that of arbitrage, that is, in its simplest form, the simultaneous purchase and sale of the same asset in two markets at different prices (S. Ross, 1976, and S. Ross, 1978): in an efficient market, the asset is correctly priced in each trading venue, and arbitrage opportunities do not exist. Should they arise, according to theory they would immediately be wiped out by arbitrageurs ready to exploit such risk-free opportunities of profit. In this sense, arbitrageurs perform the important task of 'information carriers' across exchanges. However, evidence suggests that bitcoin prices do diverge across markets, reflecting local information asymmetries (Urquhart, 2016; Cheah et al., 2018). Furthermore, a recent stream of literature shows that arbitrage in practice is de facto limited in many circumstances (Shleifer and Vishny, 1997), and recognizes the importance of better understandig who are abitrageurs in practice (Gromb and Vayanos, 2010). Identifying empirically the presence of arbitrageurs who exploit price deviations across trading venues would be a proof of its efficiency, and describing their individual behaviour at the micro level could help to understand and possibly reduce the gap that persists between theory and practice: to the best of my knowledge, this aspect is still unexplored in the financial literature at the user-specific level.

This analysis goes precisely into that direction. I seek to provide evidence of the existence of arbitrage across the Bitcoin exchange platforms, and I investigate to what extent such mechanism, that ensures the market efficiency in traditional financial markets, is conducted also in the nascent and less regulated context of the Bitcoin ecosystem. I exploit the Mt. Gox leaked dataset that contains all the internal trades executed from April 2011 to November 2013 in the Mt. Gox exchange platform and focus on the investors that performed arbitrage *across* Mt. Gox and two other markets of the Bitcoin ecosystem, namely BTC-e and Bitstamp. Most importantly, I tackle the topic through a bottom-up approach and try to identify cross-exchange arbitrage activity at the individual level: the novelty of the approach resides in the fact that the leaked dataset allows to exploit the user identifiers associated to each leg of a trade to reconstruct the *sequence* of trades performed by each investor. The methodology consists essentially in a double comparison of the user-specific sequences with the price deviations across exchanges and with the logs of anonymized trades conducted in BTC-e and Bitstamp, in order to quantify to what extent each users' trading sequences were consistent with arbitrage activity.

In this sense, this work must be intended as the conceptual extension of the analyses conducted in Chapter 3, that focus on the description of the triangular arbitrage activity *within* the same exchange platform, that is, Mt. Gox. The reported findings suggest that from 2011 to early 2013 triangular arbitrage in Mt. Gox was indeed conducted; in accordance with the anecdotal evidence, the identified arbitrageurs are few and expert users that perform many trades, conduct complex strategies - likely executing metaorders and exploiting automated trading - and whose actions are significantly more profitable with respect to those performed by non experts. I ground this work on these findings and investigate if and how price information propagates across exchanges through the activity of investors looking for mispricings in the Bitcoin markets; once identified, I describe the arbitrageurs' main features.

This analysis brings a methodological contribution by introducing new algorithms for the analysis of financial markets. To the best of my knowledge, no other study developed tools to identify individual arbitrageurs trading across platforms. It is important to note that this method is based on the contingent structure of the available data: while the user identifiers for the trades in Mt. Gox are known, the datasets of the BTC-e and Bitstamp exchanges are ledgers of anonymized trades. I thus contribute to an unexplored field by proposing a methodology valid for this specific data structure; other algorithms would be required if user identifiers were available for both exchanges, and further differences would arise in case it was possible to link the identifiers across exchanges (if user-specific identifiers are not available for any of the exchanges in exam, the only possible strategy is to analyse aggregate data).

If the proposed methodology is accepted, this work contributes also to the study of the financial markets by shedding some light on the individual behavior of users involved in arbitrage. Understanding the insights of practical arbitrage in contrast with the theoretical assumptions is a not fully answered - though relevant - question (Gromb and Vayanos, 2010); to the best of my knowledge, this work is the first describing the trading behavior of arbitrageurs between exchanges at the individual level. I find that arbitrageurs are very few users, are likely sophisticated, perform more trades and move more volume than the average user.

I also provide a substantial contribution to the growing area of research on the cryptocurrency markets, and more in general on cryptoeconomics by providing evidence that - to some extent - the mechanisms that ensure market efficiency in traditional financial markets are put into practice also in cryptocurrency markets. Second, I shed some light on the internal dynamics of such markets, and especially of the Mt. Gox platform, who played a prominent role at the early stages of the Bitcoin ecosystem. Finally, I document the history of fees paid in the three studied exchanges from 2011 to 2013.

Ultimately, hereby I extend and complete the findings reported in Chapter 3. The analysis on triangular arbitrage exploits a richer dataset (all relevant information relates to trades executed within Mt. Gox, thus arbitrage actions can be identified exactly), and evidence is stronger and more robust with respect to this analysis. However, the results obtained on two-point arbitrage refer to the role of arbitrageurs as information carriers across exchanges, and are in some sense more general.

This Chapter mainly documents a data science approach to the problem. On a high level, the pipeline consists of three stages: pre-processing, identification of potential arbitrageurs, and elimination of false positives. Each stage poses specific challenges, which I document and propose solutions for. In Section 4.2 I provide some brief context, and describe the data and the analytical approach. The most interesting parts of the work concern the definition of the metrics and the post-filtering stage (Section 4.3): first, I discuss the matching algorithm, that compares the sequences of users' actions with the price deviations across exchanges, along with the resulting similarity metrics (Section 4.3.1); then, I define a second indicator based on the identification of trades executed in a small neighborhood of time and volume (mirror trades) into the counterpart exchanges, which I call co-execution metric. Finally, I implement the post-filtering procedure by cross-comparing the two metrics, to remove the false positives (Section 4.3.2). I report the results on the set of identified arbitrageurs for the two counterpart exchanges in Section 4.3.3. In Section 4.4 I comment the results. Additional information is reported in the Appendix Section.

4.2 Background and analytical approach

4.2.1 The early stages of the Bitcoin ecosystem

The adoption of Bitcoin at its early stages[1] was mainly driven by the interest towards decentralization, enhanced privacy, and for its 'cypherpunk' libertarian values (De Filippi, 2014, and 2018). The anonymity provided by Bitcoin, combined with its digital nature, made of it the preferred method of payment for illicit activities on online black markets: an emblematic example is the case of Silk Road, an online marketplace dedicated to the commerce of illegal goods and services (Christin, 2013; Foley, Karlsen, and Putniņš, 2019). Besides this, Bitcoin has been largely associated also to other illicit activities such as thefts, frauds, and money laundering (Möser, Böhme, and Breuker, 2013; Vasek and Moore, 2015; Yin and Vatrapu, 2017). The study by Foley, Karlsen, and Putniņš, 2019 reveals that, up to 2017, around 46% of all the bitcoin transactions were associated to illicit activity, and that a significant fraction of its value can be attributed to the involvement in illicit activities.

Bitcoin at its early stages was widely used also for gambling (Möser and Böhme, 2015): the most notable example is the one of SatoshiDice, a website allowing players to bet on the result of the virtual launch of a dice with a bitcoin-based rewarding mechanism. Interestingly, the number of transactions on the network associated to gambling amount to almost 48% in the first four years of existence of Bitcoin, and are very small in volume (Lischke and Fabian, 2016). According to the authors, the largest transactions in volume are instead associated to the exchange services. Indeed, the last relevant use case of Bitcoin in its first years was as an investment asset traded in the exchange services: the most relevant in terms of liquidity and adoption was Mt. Gox, who started its activities in 2010 and dominated the market until 2013, before being filed for bankruptcy in February 2014 (Moore, Christin, and Szurdi, 2018).

[1] There is no precise classification for the different epochs of the Bitcoin ecosystem; however, the literature commonly refers to the early stages as the time window that precedes the failure of the exchange platform Mt. Gox, due to the importance that this event had within the Bitcoin ecosystem. See, e.g., Figure 2.

4.2.2 The Bitcoin exchange services

The Bitcoin design originally envisages two primary ways to obtain bitcoins: by conducting the mining activity, or by accepting bitcoins as a method of payment (Antonopoulos, 2017). In both cases, users must be active nodes of the network and possess sufficient technological skills to manage and run a dedicated software. Furthermore, the protocol itself exposes users to high risks, if mishandled (Conti et al., 2018), and the network itself can be subject to attacks and malicious activity (Benjamin Johnson et al., 2014; Vasek and Moore, 2015). Handling bitcoins correctly is non-trivial and several barriers to entry exist: especially at the early stages, the adoption of Bitcoin was confined to a limited number of users with a good knowledge and understanding of the Bitcoin design principles and mechanisms.

Currency exchanges are an alternative method to get hold of bitcoins that reduce the barriers to entry by limiting the technological and switching costs: they are trading platforms that provide users with an online interface to trade bitcoins (and other cryptocurrencies) against fiat currencies (Böhme et al., 2015). Most of them adopt the trading mechanisms of the traditional equity markets, where traders submit buy and sell orders and the exchange clears trades, providing the service for a fee. In general, they are organized as order-driven markets with two sided continuous auctions; users do not need to resort to brokerage services, and rather trade directly among them.

The most prominent Bitcoin trading platforms for the epoch I consider are Mt. Gox, Bitsamp, BTC-e, and BTC China. While the latter was active mostly on the Chinese market (and for this reason is out of the scope of this work), the first two entered the market allowing to trade Bitcoins for USD in mid-2011. Figure 42 in Appendix B.1 shows the daily number of trades executed in the three exchanges taken in exam, and for the whole time window Mt. Gox was by far the largest exchange.

A key feature characterizes such platforms: in the original design, a transaction is conducted across owners of a Bitcoin wallet (i.e., nodes of the Bitcoin network) and stored on the public ledger (Decker and Wattenhofer, 2013). Exchanges platforms instead are in all respects centralized services, as they hold in escrow the customers' funds, maintain and update the balances on a private ledger, and charge the trades with fees (in this regard such entities operate as a traditional trading platform, as the investors buy and sell an asset through an intermediary that maintains the online structure). An important consequence is that the trading operations are internal to the node and not transcribed on the blockchain, i.e., are executed off-chain (Anderson et al., 2018). To perform arbitrage across exchange platforms, in the presence of price differences, an investor could perform two simultaneous operations of withdrawal and deposit on the Bitcoin network, from different exchanges, to balance the position and make profit without risk; this operation is costly in terms of time (on average, a block is added every ten minutes, and it is good practice to wait for a few blocks to be added before considering verified the transaction), and exposes the trader to risks common to all arbitrage activity and specific to Bitcoin's design. The optimal strategy is instead to maintain a stock of both bitcoins and fiat money in accounts at each exchange, in order to react quickly to price differences (Petrov and Schufla, 2013). As a consequence, arbitrage is not observable on the public ledger; rather, evidence is embedded in the private ledgers held by the exchanges.

Notably, similar issues to those described in Subsection 4.2.1, as well as new and unconventional risks, pertain to the exchange platforms: exit scams (i.e., platforms that cease operations abruptly without refunding customers: see Moore and Christin, 2013; Vasek and Moore, 2015), price manipulation (Gandal, Hamrick, et al., 2018) and trading strategies otherwise illegal in traditional financial markets, hacks and DDoS attacks (Feder, Gandal, Hamrick, and Moore, 2018) are common phenomena to the Bitcoin exchange platforms. A report by the Bitwise company (Hougan et al., 2019) shows that several exchange platforms apparently report suspiciously inflated volumes, and according to their findings likely only few exchanges have actual volume and operate transparently (interestingly, they claim that once the fraudulent platforms are filtered out, the BTC market is smaller, more ordered and more regulated

Figure 27: Related work on two-point arbitrage.

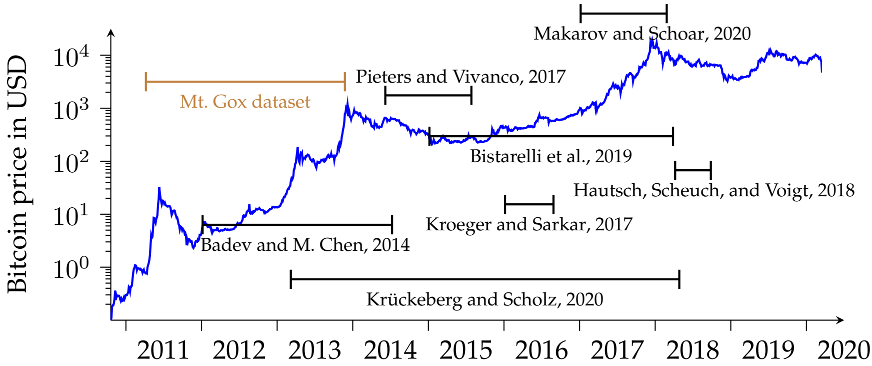

than is commonly thought). Showing that arbitrage across exchanges exists would be of the utmost importance, even more so in the light of the aforementioned challenges and risks associated to cryptocurrencies exchange platforms.

4.2.3 Market efficiency and arbitrage across bitcoin exchange platforms

A stream of literature investigates the efficient market hypothesis on Bitcoin. Urquhart (2016) studies the returns on Bitcoin price by using aggregate data from several exchange platforms from August 2010 to July 2016, and concludes that it is not efficient by performing five different tests. Nonetheless, he shows that the efficiency increases from 2013 onwards. This claim is challenged by Nadarajah and Chu (2017): they find that the odd integer power of the Bitcoin returns indeed respect the EMH for the same time window. Cheah et al. (2018) study the long-memory interdependence acoross Bitcoin markets: markets are fractionally cointegrated, inefficiencies are persistent and leave profit opportunities to investors. Similar results on long-memory interdependency are found in Bariviera (2017) and Kristoufek (2018), especially for the first years of the Bitcoin existence. The findings in Tiwari et al. (2018) are instead consistent with those in Nadarajah and Chu (2017). In summary, the literature that studies the efficiency of the Bitcoin market agrees that in the epoch I investigate (that is, until 2013) the market was not efficient, while a positive trend can be identified in the subsequent epochs.

Besides this, several previous works investigated the magnitude of arbitrage across the Bitcoin exchange platforms[2]; all the analyses are conducted on aggregate information published by the exchanges, and none of them investigated the market at the micro (individual) level. Badev and M. Chen (2014) focus on a time window comparable to the one in exam in this paper. They investigate two-point arbitrage between exchanges, but the most remarkable price differences reported fall into the period when the Mt. Gox exchange collapsed, and thus are beyond my sample. The best explanation for these differences is the counter-party risk of the tumbling exchange, rather than unexploited arbitrage opportunities. Krückeberg and Scholz (2020), who study a rather long dataset, show that arbitrage opportunities across exchanges are left unexploited, and some predictable patterns can be identified; they attribute the cause to the lack of capital and sophistication, rather than to the lack of financial tools. Pieters and Vivanco (2017) explain international inconsistencies in the price of Bitcoin with the regulatory environment towards cryptocurrencies: markets with tighter regulation, which is approximated by the level of customer identification required to open an account in the exchanges, seem to charge a risk premium. Bistarelli et al. (2019) propose a theoretical model which is fitted on empirical data. They propose a model based on the identification of common risk factors across markets that allows for small deviations to exist (and further confirm their assumptions with an empirical analysis). The work by Kroeger and Sarkar (2017), albeit incomplete at the time of writing, is probably closest to this analysis. The authors show that even after accounting for both explicit (fees) and implicit (illiquidity and volatility) transaction costs,

[2]See Chapter 3 and the works cited therein for a review of the literature on triangular arbitrage (H. Dong and W. Dong, 2015; Pieters and Vivanco, 2015; Smith, 2016; Pichl and Kaizoji, 2017; Reynolds et al., 2018; Yu and Zhang, 2018; Hirano et al., 2018; Nan and Kaizoji, 2019) and on the Bitcoin futures (Hattori and Ishida, 2018; Shynkevich, 2020; S. Lee, El Meslmani, and Switzer, 2020).

arbitrage opportunities pertain between exchanges. In particular, bitcoins are traded at a consistently lower price on one exchange, BTC-e. Makarov and Schoar (2020) investigate the role of capital controls to explain the observed market frictions between 2017 and early 2018. Finally, Hautsch, Scheuch, and Voigt, 2018 focus on a short time window from April to September 2018. Their main contribution is to derive theoretical bounds to arbitrage, which could arise as a consequence of the settlement time needed for arbitrage strategies that require bitcoin transfers across exchanges.

4.2.4 Data and Analytical approach

Figure 28 shows the analytical pipeline of the approach in detail. It is based on two data sources: the main one, the leaked internal log of files of the Mt. Gox exchange platform, and the logs of anonymized trades conducted in the two other relevant exchanges in the EUR and USD currency market for BTC-e, and USD for Bitstamp. The latter are obtained from the Bitcoincharts.com Market API[3]. I use them in the raw format to detect the mirror trades (that is, equivalent trades in a small neighborhood of time and volume) in the counterpart exchanges, and in the aggregate form at the hour level to obtain the open high low close volume (OHLCV) series to compare the mispricings between Mt. Gox and the counterpart exchanges.

The pre-processing stage involves the data cleaning of the Mt. Gox files and is discussed in Chapter 2. As in Chapter 3, I aggregate the trades belonging to the same user that occurred within the same second to account for the matching mechanism of the Mt. Gox platform's order book clearing algorithm, and I limit the analysis to the time period preceding April 2013 for the increased difficulties faced by the users in withdrawing money from their accounts (as a consequence, bitcoins started to be traded at a large premium in Mt. Gox: this is clearly visible in Figure 29, that shows the bitcoin USD price difference between Mt. Gox and the two counterpart exchanges, as a percentage of the Mt. Gox price, at the

[3] `https://bit.ly/2H52exr`. This link and the following were all accessed on 14 October 2020.

Figure 28: Overview of the analytical approach.

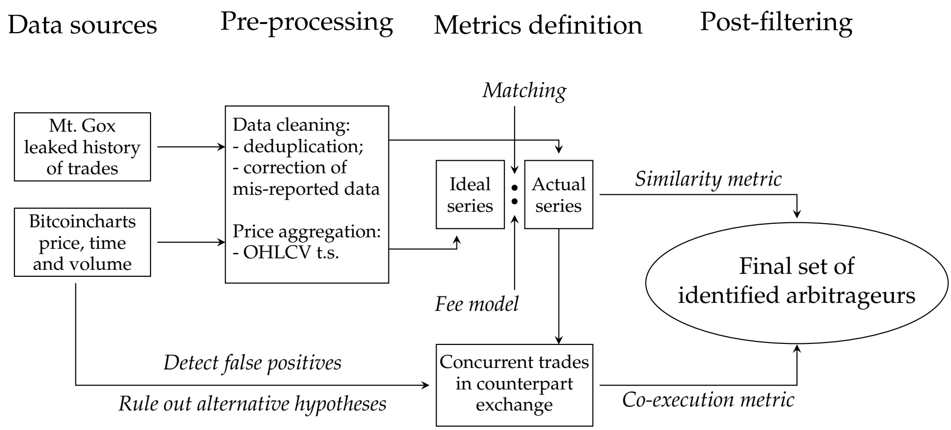

Notes: approach for the identification of arbitrage across exchanges. The pre-processing stage for the leaked dataset is common both to two-point and triangular arbitrage analysis and is described in Chapter 2.

hourly level from August 2011 to November 2013). Since I consider these events to mark the transition through different epochs, and the premium paid in Mt. Gox would likely bias the estimations, I decided to limit the analysis by excluding the trades conducted from April 2013 on. I end up considering a dataset of around 5.5 mln cleared orders, from April 2011 to the end of March 2013.

To understand the intuition underlying the metrics definition and the post-filtering stages, it is important to highlight that the internal ledger of one exchange does not allow to identify two-point arbitrage on the level of individual (pairs of) trades. However, the findings in Chapter 3 suggest that the arbitrageurs are users who perform arbitrage consistently over a longer period of time: under this assumption, it is possible to identify the users in the Mt. Gox dataset whose sequence of actions can be plausibly explained with arbitrage exploiting the differences in published prices between Mt. Gox and any counterpart exchange. The high-level idea of the approach is to generate an 'ideal' series of arbitrage actions for each counterpart of Mt. Gox from aggregate information. This ideal series is then matched against the actual trades of each user in the Mt. Gox dataset. The set of users with the highest similarity between ideal and actual trades are likely arbitrageurs. Whilst this approach is based on longitudinal similarity metric and is biased against users who execute few trades, which are excluded by the analysis, the results in Chapter 3 ensure also that the loss of information is marginal for the purposes of this analysis.

This step, alone, doesn't allow to exclude alternative explanations for the series of actions performed by those users. By definition of two-point arbitrage, an investor conducts two trades, one in Mt. Gox and one in the counterpart exchange: so, for each user and for each action I verify the existence of concurrent trades (or mirror trades) in the counterpart exchanges, and I construct a second indicator (co-execution metrics) on the share of actions with a mirror trade per each user. This approach allows to identify and remove the false positives: the post-filtering stage, that consists in cross-comparing the two metrics, further reduces the aforementioned group of potential arbitrageurs and leads eventually to the identification of the final set of arbitrageurs. The latter two stages are described in detail in the next paragraphs.

4.3 Method and Results

In this section I describe the algorithms exploited to identify the subset of arbitrageurs based on their sequence of actions. To do so, I first provide a measure of the explicit transaction costs paid by the users to trade within Mt. Gox: according to the official posted schedule, the investors would pay a fee ranging from 0.65% to 0.25% based on their individual history of volume traded. Notably, the dataset contains information on the fees paid by the user per each trade. However, I rely on the model described in Chapter 2, as inconsistencies between the fee schedule and the observed data exist[4], and I base the analysis on the estimated explicit

[4]A small fraction of legs are associated to fees higher than 1% of the trades but, most importantly, more than 1,000,000 legs pay *zero* fees. To keep a conservative approach, I

transaction costs paid by the users to perform a trade. Second, since two-point arbitrage requires *two* trades, one on Mt. Gox and a second one on another exchange, it is necessary to include an estimate of the fees that the investor would pay on the counterpart exchange. Relying on each counterpart's posted fees at the time of the analysis[5], I add a transaction cost of 0.2% of the amount traded when the counterpart exchange is BTC-e. Since Bitstamp applied a fee schedule similar to Mt. Gox, based on the volume traded in the last 30 days, but ranging from 0.5% to 0.2% and having different bands, I add a term corresponding to the fee paid in Mt. Gox, but rescaled proportionally to Bitstamp's fee range. Figure 28 in Chapter 1 reports a detailed representation of how such costs changed over time.

4.3.1 Matching and Similarity Metric

Figure 29: Price differences between Mt. Gox compared to Bitstamp and BTC-e, as a percentage of the Mt. Gox price, in USD.

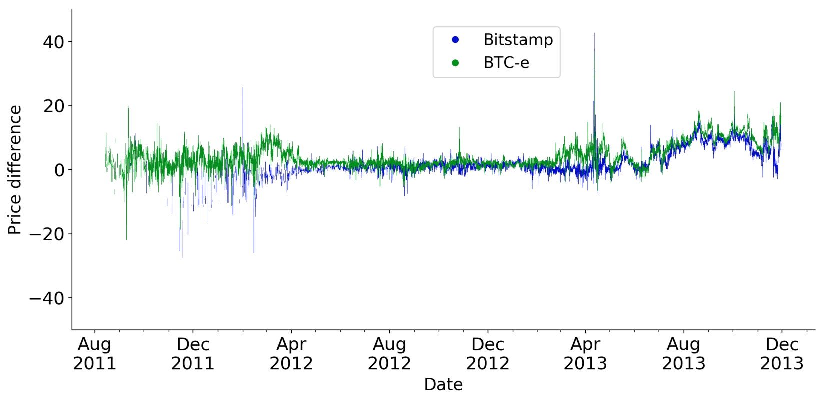

The fee estimation provides an approximation of the expected trans-

consider the model-based estimation of the fees to be more reliable.

[5]BTC-e: https://bit.ly/311ZkR9;
Bitstamp: https://bit.ly/3lPFUab.

action costs, which vary across users and time. This assumption is essential, as the individual fees affect how the same arbitrage opportunity is perceived by different users: an identical mis-pricing can appear as an unexploited arbitrage opportunity for an investor who pays low enough fees, and a costly operation for a trader that pays higher fees (it is less clear and hard to determine, instead, what are the hidden costs that the investors face). The next step of the analysis consists in defining a heuristic approach to identify the users who are likely to be arbitrageurs: I *match* the sequence of actions executed by each real user, together with the associated fees paid, against an ideal sequence of optimal actions made by a 'perfect arbitrageur'; thus, I introduce a *similarity metric* to classify investors in potential arbitrageurs and non-arbitrageurs. The comparison of the mispricings is based on hourly frequency, therefore I first aggregate the price time series for the two counterpart exchanges and compare it with Mt. Gox hourly prices, as Figure 29 shows.

The matching procedure works as follows. First, I describe the rational behavior of an ideal arbitrageur *within* Mt. Gox, with respect to each observable hourly mispricing between Mt. Gox and a counterpart exchange E. For every time interval, the optimal response of an ideal arbitrageur is drawn from the set of actions A = {Buy (B), Sell (S), Hold (H)}: without transaction costs, should the prices be higher in Mt. Gox, an active arbitrageur would sell in Mt. Gox (and buy in E). Viceversa, he/she would buy in Mt. Gox when prices are lower. If there is no mispricing, the best action is to hold (in practice, without transaction costs, this never happens). I construct then a sequence of hourly ideal response actions (B,S,H) to price changes between Mt. Gox and E. The magnitude of the mispricing allows to measure also the associated maximum transaction costs that would entail zero profits in each hour. The information contained in this ideal sequence is common to each user and publicly available.

Second, I aggregate user actions at the hour level. When a user submits more than one trade in the same hour, I define the *hourly prevalent action* as Buy if the difference between the amount of bitcoins bought and

Figure 30: Illustrative example of the matching heuristic.

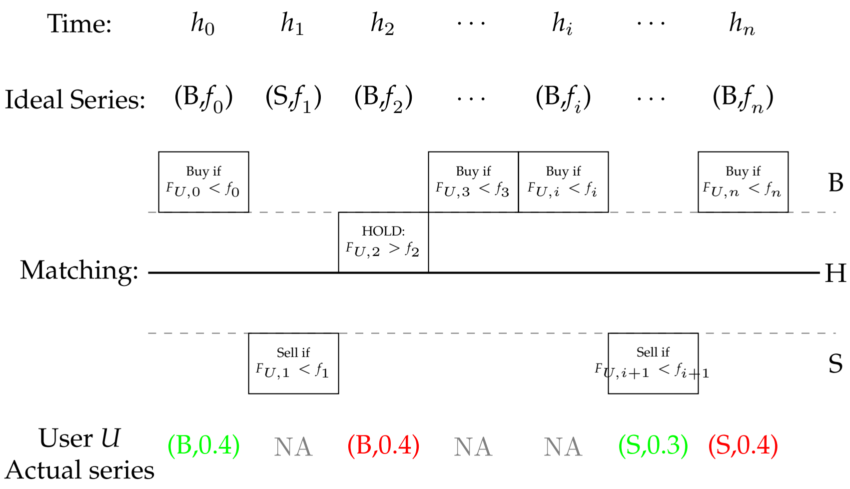

Notes: at each point in time h_t the buy/sell action that an ideal arbitrageur would execute given the mis-pricing is known and public, as well as the value f_t for the magnitude of the arbitrage opportunity. $F_{U,i}$ represents the fee paid and varies across time and users. The last row shows the tuples (Prevalent Action, Fee) for a representative user U. Now, for this illustrative example, suppose that all the values f_i are bigger than 0.5%, except $f_2 = 0.3\%$. Once individual fees are included, some of the f values might turn out to be too small for the user to be exploited as arbitrage opportunities, because the fees exceed the size of the mis-pricing. E.g. in this case, even though the ideal and actual action at time h_2 correspond, the fee exceeds the price differences: an ideal arbitrageur would then 'Hold' (H); so, this action is counted as 'non-arbitrageur like' for user U. NA stands for Not Active, meaning that in this time window the user did not make any trade; green tuples are classified as 'arbitrageurs-like', while red tuples as 'non arbitrageurs-like'.

sold is positive, and as Sell if negative[6]. Matching takes place only on the hours in which the user is active[7]. I do not make any deduction on the 'arbitrageur-like' behavior of a user during times of inactivity.

[6]I exclude the (rare) cases in which the difference is equal to zero. Fees are aggregated by computing the mean value for trades within the same hour.

[7]Note also that the matching algorithm only compares time-periods in which *both* the user action and the price difference between Mt. Gox and the counterpart exchange are available. E.g., Bitstamp data are available only from September 2011 onwards.

Third, I compare the hourly series of actions of each user to the ideal one, updated so as to account for his/her individual expected fees. The intuition is that, to execute arbitrage, the user observes in each time interval the publicly available information on price deviations and maximum transaction costs, and then he/she incorporates the private information on the fee expected to pay given the personal history of volume traded. Thus, given the *user-specific expected fees*, he/she derives the *individual* ideal best action: once such costs are accounted for, the same arbitrage opportunities appear profitable (or unprofitable) to a different degree for each user.

The matching algorithm compares the user sequence to his/her individual ideal sequence of actions: an action is considered as 'arbitrage-like' if it corresponds to the one that an ideal arbitrageur would perform in the same time window, and if concurrently the user-specific expected transaction costs do not exceed the size of the mis-pricing. If the expected transaction costs are too large, then the action would be not profitable. Figure 30 shows an illustrative example of the matching procedure for an illustrative user U.

As a result of the matching algorithm, only a fraction of each user 'hourly actions' correspond to potential two-point arbitrage actions. For each user I compute then the fraction of 'arbitrage-like' actions that correspond to potential two-point arbitrage with respect to the counterpart exchange E. I rank the users based on this parameter in descending order. I repeat these steps for each counterpart exchange, considering the price time series denominated in USD. Figure 31 shows two lines, one per exchange, each dot representing the users ranked in descending order, based on the share of 'arbitrage-like' actions. Note that, for each counterpart exchange, actions are matched independently: a user ranked n^{th} in a counterpart exchange might be ranked m^{th} in another, and a user can be classified as potential arbitrageur in no exchange as well as in both exchanges. As stated in Section 4.2.4, I assume that arbitrageurs are users who perform arbitrage consistently over a long period of time; thus, for each counterpart exchange, I take into account only the users for which the matching occurs on more than 100 hours. Users

Figure 31: Share of trades that correspond to two-point arbitrage actions with another exchange, per user.

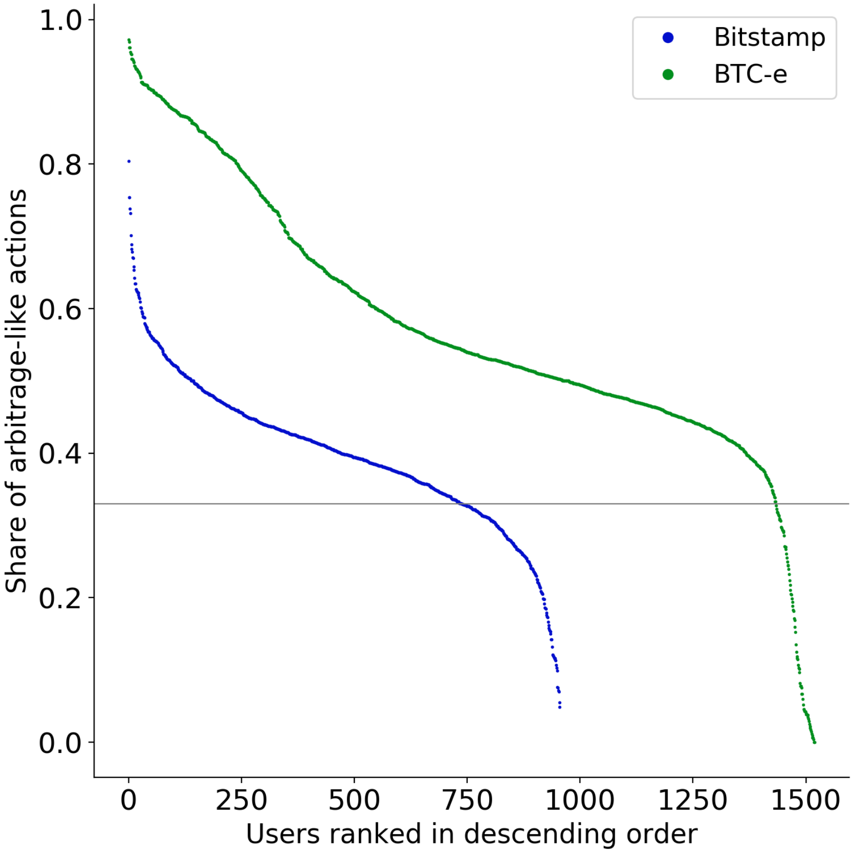

Notes: trades are evaluated independently for the counterpart exchanges Bitstamp and BTC-e; users whose share is above 0.33 are treated as potential arbitrageurs. The analysis is performed for users for which the matching occurs on more than 100 hours.

whose share of 'arbitrage-like' actions is above 0.33 are classified as potential arbitrageurs (see Appendix B.1 for further details on the choice of the parameters). On a total of 71,808 users, 70,367 are always classified as non-arbitrageurs, while 1,441 investors are detected as potential arbitrageurs on at least one exchange. The set of users identified on Bitstamp is smaller with respect to BTC-e (around 750 potential arbitrageurs against around 1400): the dominance of BTC-e is due to the fact that, throughout the whole time window, bitcoins are traded at a slightly lower price in BTC-e. Both lines follow a sigmoid trend; in Bitstamp, with respect to BTC-e, it is possible to notice a steeper decline of the users whose fraction of arbitrage-like actions is high.

4.3.2 Post-filtering and results

The potential arbitrageurs identified with this method are likely an overestimation of the true set of investors that performed arbitrage; most importantly, the algorithm just described relies on assumptions which imply that their sequences of actions *could* be explained with arbitrage, but it doesn't allow to rule out the possible alternative hypotheses on the detected users' trading behavior. I implement a second procedure to overcome these issues and to control for the presence of false positives; as a result, I further narrow down the set of identified potential arbitrageurs.

By definition of two-point arbitrage, whenever an arbitrageur pursues his/her strategy, one would expect to observe two trades, executed by the same investor, to buy (sell) bitcoins in Mt. Gox and simultaneously sell (buy) the same amount in the counterpart exchange. However, while the Mt. Gox trades are labelled at the user level, the logs of the counterpart exchanges are anonymized. Thus, I cannot identify exactly the arbitrage actions across exchanges by comparing the user identifiers, and the absence of user identifiers implies also that is not possible to infer the direction of the action (whether buy or sell) executed in the counterpart exchange. The most informative approach is thus to verify for which trades a concurrent trade in the counterpart exchange does exist: exploiting again the assumption that arbitrageurs are investors that performed

Figure 32: Identification of the mirror trades.

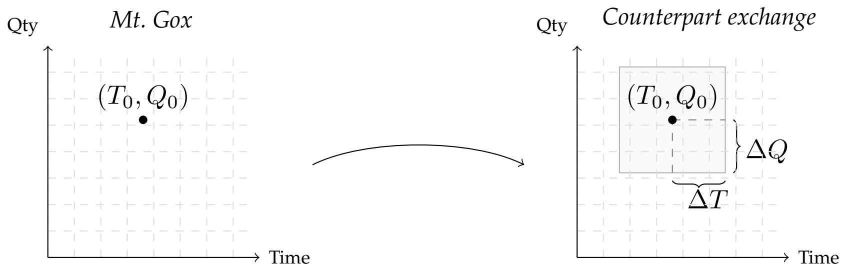

Notes: for each trade executed at time T_0 and volume Q_0 by user U in Mt. Gox (left), I verify whether there exists or not a trade in a neighborhood $\pm\Delta T$ and $\pm\Delta Q$ in the counterpart exchange E (right). They gray shaded area represents the region of values [T,Q] accepted to count a trade in E as a mirror trade for (T_0, Q_0) in Mt. Gox. To identify exactly an arbitrage action, additional information on the orders executed in the counterpart exchanges would be needed: as their logs are anonymized, I do not have access to the user identifier; additionally, the direction (whether a buy or sell) of the order in the counterpart exchange is not known.

arbitrage for long periods of time, I can construct a second indicator that measures, for each user, the share of actions in Mt. Gox with a mirror trade on the counterpart exchange. This co-execution metric gives additional information on the probability that an investor is an arbitrageur. Furthermore, it obviates a drawback of the similarity metric introduced before, which compares price differences at hourly level, while this step is more oriented towards a high-frequency approach. Note that, while one would expect to find a high number of mirror actions for the users identified as potential arbitrageurs, there's no reason to think that there should be any kind of relationship, for the non-arbitrageur users, among the execution timings of trades in Mt. Gox and in the counterpart exchanges.

I consider again only the set of users whose sequence of trades match for at least 100 time intervals. For each of their trades I check for the presence of a *mirror* trade in the counterpart exchange, that is, a trade in a reasonably small neighborhood of time and volume, $\pm\Delta T$ and $\pm\Delta Q$. Figure 32 shows the intuition behind this step. Consistently with the

Figure 33: Share of actions with a mirror trade in the counterpart exchange.

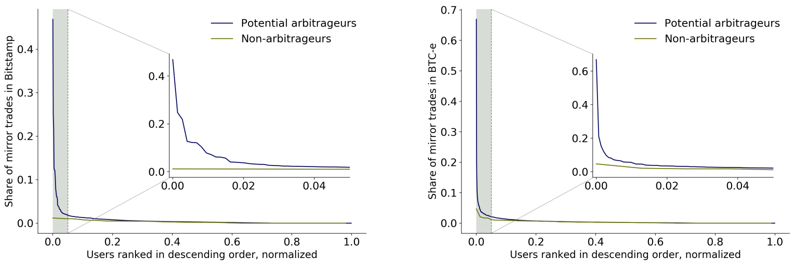

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

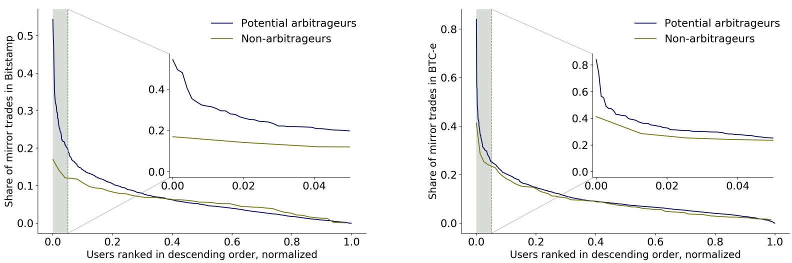

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Notes: left: Bitstamp, right: BTC-e. Users ranked in descending order. The matching occurred on at least 100 hours for all the users reported in the plots. Panel (a) reports the results for the more restrictive scenario, panel (b) for the less restrictive one. Users are represented on two different lines (yellow: non-arbitrageurs, blue: potential arbitrageurs), to qualitatively inspect the differences between the two groups. The x-axis spans from 0 to 1 (the number of users in the two groups is normalized, so that distributions are easier to compare): each plot includes a subplot that focuses on the 5% of arbitrageurs and non-arbitrageurs with the highest share of mirror trades.

Figure 34: Post-filtering procedure.

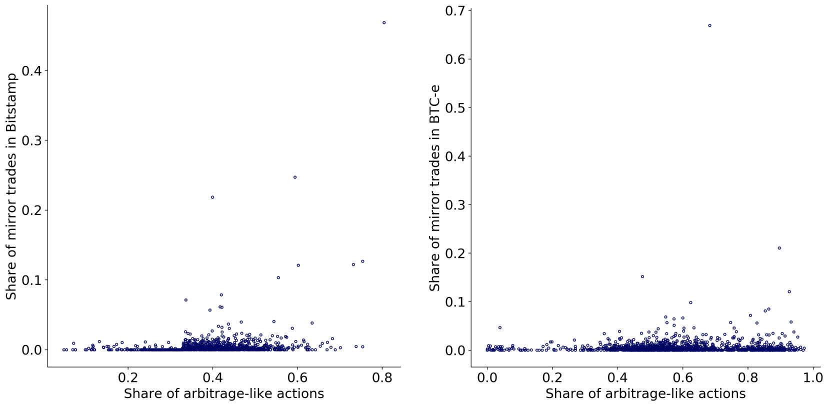

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

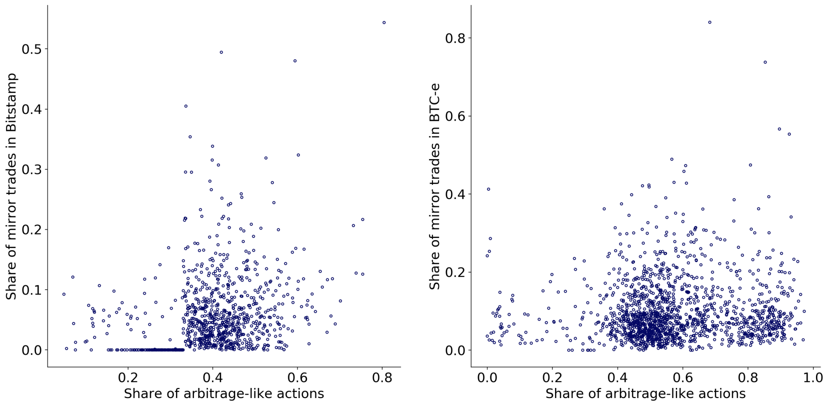

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Notes: I plot a dot per each user for which the matching occurs on more than 100 hours. The x-axis is the share of arbitrage-like actions, and the y-axis is the share of actions with a mirror trade in the counterpart exchange (left: Bitstamp, right: BTC-e). Again, Panel (a) reports the results for the more restrictive scenario, and Panel (b) for the looser one.

empirical evidence reported in Chapter 3, I consider two different cases: a more restrictive scenario with $\Delta T = 30s$ and $\Delta Q = 1\%$, and a less restrictive one with $\Delta T = 300s$, $\Delta Q = 10\%$.

All plots in Figure 33 report only the users for which the matching occurred on at least 100 hours. As described above, the classification as potential or non arbitrageur on one counterpart exchange is independent with respect to the classification in the other. Panel (a) reports the results for $\Delta T = 30s$ and $\Delta Q = 1\%$. Users are ranked in descending order, based on the share of actions that have a mirror trade in the counterpart exchange (Bitstamp on the left, BTC-e on the right). In each plot the blue line represents the set of potential arbitrageurs with respect to the counterpart exchange, and the yellow one shows the set of non-arbitrageur users. The number of users in the two groups is normalized (the x-axis spans from 0 to 1), so that distributions are easier to compare. Each plot contains a subplot that focuses on the 5% of users with the highest share, both for the group of arbitrageurs and for the group of non-arbitrageurs. The vast majority of the users have values equal or close to 0 in both sets. All the (few) users with a high share of mirror trades are in the set of the potential arbitrageurs. The results are similar in the two counterpart exchanges. Panel (b) repeats the same analysis in a less restrictive scenario ($\Delta T = 300s$ and $\Delta Q = 10\%$). The general results obtained in the previous case are valid also in this scenario. The users with a high share of mirror actions are few and belong to the group of potential arbitrageurs; the behavior of the users in the tail slightly differs: in both plots in Panel (b) the share of actions with a mirror trade grows up faster. It is likely that this effect is attributable to noise.

4.3.3 Identification of the arbitrageurs

To post-filter the group of potential arbitrageurs, I cross-compare the two metrics for each user in the analysis; I show the results in Figure 34, again both for $\Delta T = 30s$, $\Delta Q = 1\%$ (top) and for $\Delta T = 300s$, $\Delta Q = 10\%$ (bottom). The plots on the left concern Bitstamp, those on the right BTC-e. Each dot represents a user, defined by the share of arbitrage-like actions (x-axis) and the share of mirror actions in the counterpart exchange (y-axis). The plots show that, especially for the more restrictive scenario, the limited number of users with a high share of mirror trades also have a high share of arbitrage-like actions. In particular, two different users stand out when compared to the others, one in Bitstamp and one in BTC-e: I refer to them as $U_{Bitstamp}$ and U_{BTCe}. I describe their trading behavior in Table 20. In addition to that, I provide a description in Table 20 of a wider group of users with at least 30% of mirror actions and at least 30% of arbitrage-like actions in the less restrictive scenario, which I classify as arbitrageurs. The choice of these two thresholds is arbitrary and is discussed in Appendix B.1. The plots in Figure 34 show that the identification procedure is more sensitive to the measure of the share of mirror trades.

I identify 49 arbitrageurs in total, 10 between Mt. Gox and Bitstamp, and 45 between Mt. Gox and BTC-e. Six of them are detected as arbitrageurs against both Bitstamp and BTC-e. Table 20 describes the trading behaviors of the users detected as arbitrageurs. Column (e), that reports the descriptive statistics for the users active for more than 100 hours, is a benchmark for the comparison of $U_{Bitstamp}$ (column (a)), U_{BTCe} (column (b)), and the users with at least 30% of mirror actions and at least 30% of arbitrage-like actions with respect to BTC-e and Bitstamp (respectively columns (c) and (d)). Column (f) shows the results for the whole sample of users. $U_{Bitstamp}$ and U_{BTCe} are more active in terms of number of trades, Bitcoins, and fiat money traded, when compared to the full sample of users (column (f)), while they are comparable to the set of users described by column (e). The difference with respect to the set in (f) can be explained as the analysis structurally excludes the users with few transactions. It is less trivial instead to determine the reasons why the differences with group (e) are smaller.

The second relevant finding is that only a small group of users are identified as arbitrageurs. Both elements are in accordance with the assertion that arbitrageurs are few and expert investors 'with the knowledge and information to engage in arbitrage' (Shleifer and Vishny, 1997). Focusing on the two groups in columns (c) and (d), the first noticeable pattern is that the characteristics of those in BTC-e, column (c), and Bitstamp, column (d), are quite different, with the latter being composed of users more active both in terms of quantity of trades performed and of volumes moved. The identified arbitrageurs in Bitstamp and BTC-e performed on average less trades (both Buy and Sell) and moved less money with respect to users in (e). Interestingly, all the users described by (c), (d) and (e) tended to be more sellers than buyers of Bitcoins. Unexpectedly, both groups of arbitrageurs alternated buy and sell actions less often than group (e); the user identifier of (e) is higher, suggesting that these users on average entered the market earlier with respect to the users described in the other columns. While the behavior of the identified arbitrageurs is definitely different with respect to that of the whole sample of users (column (f)), it is less straightforward to find a trading pattern that differentiates them from the other investors active for long periods of time (column (e)).

Table 20: Comparison of identified arbitrageurs to non-arbitrageurs

		Users with > 100 hours of activity					All users
		Identified arbitrageurs				All	
N of users		U_{BTC-e}	$U_{Bitstamp}$	BTC-e (45)	Bitstamp (10)	$(2,039)$	$(71,808)$
		(a)	(b)	(c)	(d)	(e)	(f)
Total trades	mean	1,203	5,970	942.8	1,458.3	1,908.5	77.8
	50%			549	515.5	523	7
USD trades	mean	1,203	5,970	940.6	1,458.3	1,827.2	70.6
	50%			549	515.5	513	4
Buy trades	mean	381	1,596	365.5	683.0	896.5	38.4
	50%			188	305	222	4
Sell trades	mean	822	4,374	577.3	775.3	1,012	39.4
	50%			317	211	288	1
Changes of State	mean	84	305	202.9	136.9	373.7	15.2
	50%			103	76.5	109	0
Bitcoins bought	mean	154.7	4,101.6	936.3	2,967.7	12,031.2	533
	50%			258	1,386.3	1,425.5	18.2
Bitcoins sold	mean	260.7	18,974.8	1,417.4	4,269.4	13,160	533
	50%			420.4	1,254.2	1,869	2
'\$' sent	mean	8,550	46,728.6	31,171.0	68,245.6	121,683.7	6,000
	50%			6,202.4	40,165.2	13,593.6	237.7
'\$' received	mean	14,292.1	231,780.8	37,717.8	81,536.3	130,316.9	6,000
	50%			11,170.9	46,029.4	18,285.7	31.3
'\$' per trade	mean	19	46.7	73.9	185.0	149.5	158.9
	50%			34.5	72.4	66.8	67.1
Active hours	mean	167	1,307	261.8	324.4	387.1	20.7
	50%			211	131	188	4
Active days	mean	45	309	83.0	85.5	119.9	10.2
	50%			71	50.5	92	3
Trades per hour	mean	7.2	4.6	3.2	3.6	3.3	2.2
	50%			2.4	3.4	2.5	1.5
Trades per day	mean	26.7	19.3	14.5	16.7	12.7	3.2
	50%			8.7	11.0	6.1	2
User ID	mean	52,34X	31,20X	35,829.0	32,306.1	19,593.1	35,903.5
	50%			3559X	33,40X	18,11X	35,90X

Notes: A note on the meaning of '\$'. Some users in columns (e) and (f) made transactions in fiat currency other than USD. Thus, to make results comparable, I converted in USD the value of the trades denominated in different fiat currencies. Note also that the user ID is a sequential value that does not correspond to the one originally reported in the leaked dataset.

4.4 Discussion

The aim of this work is to mine Mt. Gox's internal ledger in order to identify users who likely carried out two-point arbitrage between Mt. Gox and two other Bitcoin exchanges, Bitstamp and BTC-e. To do so, I pre-process the data sources following related work. I then exploit the fee model introduced in Chapter 1 to estimate user-specific transaction costs, and I define a similarity metric to compare users' actual trades to the actions of an 'ideal arbitrageur', given the individual transaction costs and the price differences between Mt. Gox and the two counterpart exchanges. I introduce a second indicator that measures, for each user, how many trades in Mt. Gox have a mirror trade (that is, a trade in a small neighborhood of volume and time), on a counterpart exchange. I leverage these two metrics to identify a set of users that I classify as arbitrageurs between Mt. Gox and the counterpart exchanges.

The paper contributes to the literature on Bitcoin and cryptocurrencies by shedding light on the internal dynamics of the Mt. Gox ecosystem, and in particular on the actual fee scheme paid by the users in Mt. Gox (and in the counterpart exchanges) from April 2011 to November 2013. Most of all, I devise a heuristic approach to detect arbitrageurs across exchanges and I describe their trading behavior, giving a contribution on the financial literature related to the identification of arbitrageurs: to the best of my knowledge, there is no empirical study investigating the arbitrage activity *across* trading venues through a micro level approach. As far as I know, this work represents the first attempt to fill this gap in the literature, and the empirical evidence in support of the thesis that arbitrageurs are 'few professional, highly specialized investors' (Shleifer and Vishny, 1997, p. 36), is limited thus far to the contribution in Chapter 3.

Assuming that the devised strategy actually detects correctly the arbitrageurs, a further extension of the work would consist in quantitatively determining the impact of the arbitrageurs' actions on the spreads between prices in Mt. Gox and in the counterpart exchanges. As Figure 29 shows, after a first phase characterized by high volatility, prices tended to converge from around April 2012 to around March 2013. A relevant related research question would be to investigate whether the detected users had a role in reducing the spread during this epoch and in achieving higher market efficiency.

Some open problems remain to be solved. For instance, I am aware that the proposed method needs to be validated and requires further robustness checks: the heuristics rely on the choice of partially arbitrary parameters which are not fully supported by formal statistical methods nor by theoretical argumentations. Besides this, a general limitation of this approach and dataset is that it is not possible to observe the individual behavior on the counterpart exchanges. Only the access to this information, or an explicit proof from the same users that performed arbitrage, would strengthen the evidence that what is being identified is actually two-point arbitrage. As a consequence, rather than individual arbitrage actions, I identify trading strategies plausibly consistent with arbitrage, introducing a bias towards users who conducted *mostly* arbitrage in their trading activity, and penalizing those who concurrently conducted arbitrage *and* other strategies. E.g., a user who conducted arbitrage only for a limited time window before moving to other strategies may not be classified as arbitrageur. This bias is caused by the introduction of a cutoff, which is however necessary to reduce the number of false positives with low scores for the devised metrics. Furthermore, this longitudinal similarity metric is biased against users with few transactions: this restricts the ability to comprehensively answer the question if the arbitrageurs are a few big players or many small investors. This research question is better addressed in Chapter 3. However, as discussed above, the empirical evidence reported therein shows that the vast majority of triangular arbitrage actions are executed by few large users. The loss of information due to the exclusion of users who executed few trades thus seems to be marginal.

Another drawback is that it is not trivial to interpret arbitrage with more than one counterpart exchange involved for the same user because an investor could engage in two-point arbitrage between pairs of exchanges not including Mt. Gox. In this regard, a possible way forward could be to compose ideal series with actions involving two-point arbitrage between all pairs of exchanges and consider actions that do not involve Mt. Gox as censored data points. However, in the time window considered Mt. Gox was by far the largest platform in terms of volumes traded: it is likely that all arbitrage activity had to be executed through Mt. Gox.

Even when all of these open problems are addressed, the external validity remains a concern. Can one learn anything substantial about arbitrage activity across trading venues in conventional financial markets from a nascent niche market that requires technical sophistication and willingness to accept unusual risks? This question must be asked for the early years of Bitcoin as well as the emerging markets in the crypto-token ecosystem. For example, a very recent related work studying automated arbitrage on Ethereum's decentralized token exchanges reports sizable opportunities, which are routinely exploited by a range of competing trading bots (Daian et al., 2019). However, the external validity is further compromised by an observer effect: apparently, the race was triggered by a blog post and the release of proof-of-concept code for a trading bot by members of that research team!

Chapter 5

Conclusions

In the previous Chapters I described the methodologies that I devised to identify the investors that conducted triangular and two-point arbitrage within the Mt. Gox exchange platform between April 2011 and March 2013.

The main contribution of my thesis is that it is the first comprehensive empirical study of the arbitrageurs at the individual level: to the best of my knowledge, the previous empirical academic literature either provided studies at the aggregate level or anecdotal evidence regarding the arbitrageurs at the micro level; in addition, the theoretical description of the arbitrageurs in traditional finance is far from the one provided by the practitioners, and many recent studies show that de facto limits to arbitrage exist, taking the form of persistent mispricings. Shedding some light on the true behaviour and trading patterns of the arbitrageurs is a compelling issue for the current research in behavioural finance (Gromb and Vayanos, 2010). In the best case scenario, answering extensively to the question 'who are the arbitrageurs?' could help refine the assumptions for the theoretical studies that derive optimal trading strategies in the presence of arbitrageurs; a better understanding of the underlying decision-making processes at the individual level could become a key factor to further investigate the causes of the limits to arbitrage. In this case, such investigations could have significant impact also on real markets, leading to an overall reduction of the market inefficiencies.

In this sense, my study is a pioneering work in this stream of literature, as it identifies for the first time the users who conducted arbitrage in a financial market, and it provides empirical evidence of their trade patterns. In contrast with the assumptions typical of the traditional theoretical finance, and coherently with more recent studies in behavioural finance (e.g. Shleifer and Vishny, 1997), the main finding of my work is that the arbitrageurs are indeed few and skilled investors.

Nonetheless, I am aware of the potential limitations of this analysis, as it is not trivial to prove the validity of these results in other financial markets. Such external validation is made even more complex by the specificity of the empirical setting: I focus on the Bitcoin market, which is poorly regulated, subject to high unconventional risks, and I consider a time period that corresponds to the early stages of the Bitcoin adoption. This choice is opportunistic, in the sense that I exploit the Mt. Gox leaked dataset for its unique richness: the trades are recorded in the private ledger and are labelled by user identifiers, making possible to reconstruct the sequences of individual actions, and exploit this essential feature to identify the arbitrage activity.

I envisaged and devised two different approaches to identify the two main types of arbitrage conducted in cryptocurrency markets, that is, arbitrage within a single cryptocurrency market (triangular), and arbitrage across multiple currency markets (two-point). Each method poses specific challenges, which I document and propose solutions for.

Chapter 2 provides a twofold contribution to the field of cryptoeconomics by enriching the literature on the cryptocurrency exchange platforms. First, I supply a polished version of the Mt. Gox leaked dataset. I preprocess it by removing duplicate rows, and I correct misreported entries; then I merge it with additional information published by the Mt. Gox website itself. I provide a data cleaning procedure that relies on and improves upon the existing ones, and I show that my polished dataset is consistent with the ones obtained by Scaillet, Treccani, and Trevisan, 2017, Feder, Gandal, Hamrick, and Moore, 2018, and Gandal, Hamrick, et al., 2018. The comparison with the data provided by Bitcoincharts.com confirms the validity of the proposed deduplication technique. Second, I provide and discuss a model to account for the explicit transaction costs borne by the users. I show that the fees a user would expect to pay given their history of volume traded are consistent with the official fee schedule reported by the exchange platform, and I provide possible explanations for the trades that are associated with zero fees. Indeed, most of them are executed by few users or are executed in specific time windows. Besides this, I provide a clear scheme for the fees paid by the users in the most relevant exchanges active in the Bitcoin ecosystem from 2011 to 2013 (Mt. Gox, Bitstamp, BTC-e).

The most relevant results are reported in Chapter 3, where I describe the methodology used to identify the triangular arbitrage executed within the Mt. Gox platform. The exceptional granularity of the leaked dataset allows to identify *exactly* the triangular arbitrage actions: indeed, users were allowed to trade within the same platform in multiple fiat-to-bitcoin markets, and were entitled to have only one personal account at a time. This, along with the availability of user labels associated to the legs of all trades, enables to match the pairs of (buy,sell) legs that satisfy the textbook properties of arbitrage: that is, two legs executed by the same investor, executed in two separate trades (different trade ID), using different currencies, and such that the time delay, δT, and the volume difference, δQ, are smaller than (or equal) the maximum boundaries ΔT and ΔQ.

Thus, the triangular arbitrage activity can be observed completely within the private dataset of one exchange alone. By imposing a conservative boundary ($\Delta T = 300s$ and $\Delta Q = 10\%$), I identify N = 6,629 actions, executed by N = 440 users. Most of the actions have values of $\delta T \approx 0s$ and $\delta Q \approx 0\%$. This does not come unexpectedly, and it matches the textbook definition according to which arbitrage is performed through simultaneous actions involving almost equivalent securities. The arbitrage activity is distributed heterogeneously across the investors who conducted it: few users are responsible for the majority of actions. Interestingly, such large users always exploit multiple fiat-to-fiat markets (i.e., they execute arbitrage using bitcoins as the vehicle currency against more than two fiat currencies); a fraction of their arbitrage actions is often executed in the form of metaorders; they only use limit orders and never trade aggressively (which lets hypothesize that they independently formed similar estimations of the price execution risk associated to market orders and of the execution risk associated to limit orders). In summary, their trade patterns are complex, and I assume that these four variables are a proxy of expertise. I define the profitability of an arbitrage action as the percentage difference between the implied exchange rate and the official rate (by construction the action is profitable if the former is larger than the latter). Indeed I show that, whatever the measure of expertise considered, the arbitrage actions executed by expert users are more profitable, with respect to those executed by non expert users, which are on average non profitable when the explicit transaction costs are accounted for. The premium is statistically significant: as an example, for the dummy variable $D(Currency)$, which is equal to 0 when the user executed arbitrage actions only in a single fiat-to-fiat market and 1 otherwise, the premium between experts and non experts amounts roughly to 1.5%. In addition, I show that when within-user (across hours and markets) variation is exploited the expert users who conduct arbitrage in multiple markets make profits by reacting quickly to plausible exogenous variations on the official exchange rates. The better ability in incorporating information from the markets on volatility and the correct calculation of the transaction costs, which ultimately result in a better timing choice of the skilled users at a small scale level, are in my interpretation the most relevant factors that explain the differences in profits between sophisticated and non sophisticated users.

The main findings of this Chapter try to provide an answer to some aspects of the important question 'who are the arbitrageurs?'. They concern the identification of the triangular arbitrage activity, the description of the arbitrageurs' characteristics and of their trade patterns, and the design of an identification strategy showing that significant differences arise across expert and non expert users. All these aspects can contribute significantly to the field of behavioural finance, and could help reducing the gap between the theoretical and practical description of arbitrage.

First of all, the devised methodology identifies a non-negligible number of actions that are coherently explainable as arbitrage activity, while it is hard to find alternative explanations for their execution. Even in the case that some of them are false positives, or that my method does not detect correctly all the actions (e.g., due to the boundaries ΔT and ΔQ), the analysis provides evidence that the triangular arbitrage was indeed conducted within Mt. Gox even at the earliest stages of the Bitcoin ecosystem. Second, I provide extensive proof that the arbitrageurs are few (only 440 out of around 72,000 are involved in arbitrage), and that the activity of the sophisticated ones is systematically more profitable. To identify the arbitrage activity, I introduced new algorithms suitable to investigate the financial markets when user identifiers are available and the individual sequences of actions can be reconstructed. Also this aspect represents a methodological contribution to the literature: to the best of my knowledge, no other study developed similar tools, and such methodology can be easily extended to any (non) financial market. The results are also relevant for the field of cryptoeconomics and of the cryptocurrency exchanges, as they bring evidence that established and relevant mechanisms that guarantee the market efficiency in traditional finance do apply even to the Bitcoin ecosystem, despite being a nascent and poorly regulated market, whose design principles are unconventional and complex. Clearly, this does not mean that the market is completely efficient and that all opportunities are exploited. This question goes beyond the scope of my work. Evaluating to what extent the existing opportunities of arbitrage are exploited could be subject of further research.

Chapter 4 extends the analysis to the arbitrage *across* exchange platforms trading bitcoins against fiat currencies (Mt. Gox, Bitstamp, BTC-e). In this case it is not possible to identify exactly the arbitrage activity, as the trades in the Mt. Gox counterpart exchanges are not labelled at the user level. Thus, I adopt a different methodology, based on comparing the individual sequence of actions of each user to the ideal sequence of actions an ideal arbitrageur would perform, given the mispricings between the two exchanges and the individual estimation of the transaction costs. I find that only around 1,400 investors are potential arbitrageurs.

It is worth noting that, unlike in Chapter 3, in this case it is harder to rule out alternative possible explanations to the observed sequences of actions; thus, I post-filter the set of potential arbitrageurs by measuring, for each potential arbitrageur, the share of actions in Mt. Gox for which it is possible to find a similar trade (small time delay and volume difference) in the counterpart exchange. This fraction will be high for an arbitrageur and low for a non-arbitrageur, and thus is intended to remove the false positives detected in the first step of the procedure. The final identification relies on the cross-comparison of the two metrics, that reduces the number of identified arbitrageurs to 10 with Bitstamp and 45 with BTC-e.

Considerations similar to those reported for the methodology devised for the triangular arbitrage apply also for the one implemented in this Chapter: the new algorithms introduced represent a methodological contribution to the literature on the quantitative analysis of the financial markets. The most important finding is that this analysis confirms the results found in Chapter 3, and it represents a contribution ascribable to the field of behavioral finance: the post-filtering stage greatly reduces the number of users classified as arbitrageurs, which are few (N = 49) and take positions larger than the average user. However, this methodology exploits less rich datasets, and the users who executed less than ten trades are excluded in the analysis by construction. This limits the possibility to comprehensively answer, in this scenario, the question if the arbitrageurs are a few big players or many small investors, as it is not possible to provide information for the latter group. Nonetheless, the findings in Chapter 3 suggest that this loss of information has a limited impact on the analysis. The second relevant aspect is that, though the evidence is less strong, this analysis is more general and implies that some investors conduct the important function of 'information carriers' across exchanges, certifying also in this case the existence of the mechanisms that ensure the efficiency in traditional markets.

A final remarkable observation stands from a comparison of the groups of arbitrageurs identified with the two methodologies: interestingly, none of the arbitrageurs who executed two-point arbitrage is identified also as triangular arbitrageur. This finding is coherent with a certain degree of specialization of the arbitrageurs, another relevant feature attributed to practical arbitrageurs. Indeed, the two strategies are relatively different and imply a cost in devising each of them. This result suggests thus that arbitrageurs focused on a single strategy.

In conclusion, the results reported in Chapter 3 and Chapter 4 are coherent and solid. While it is not possible to utterly exclude the presence of false positives in the identified groups of arbitrageurs, a last form of validation of this work could come in the form of a direct proof from the same arbitrageurs that conducted arbitrage in the Mt. Gox platform.

Appendix A

Supplementary material on triangular arbitrage

A.1 Additional robustness checks on ΔT and ΔQ

As described in Section 3.3, the algorithm implemented to identify the arbitrage actions compares the legs executed by the same user in a small neighborhood of time and volume, defined by the parameters ΔT and ΔQ. In this Appendix I show that the results of my estimations do not vary significantly when I consider larger or smaller sets of identified arbitrage actions — i.e., if I modify the boundaries for the time delay and volume difference by varying ΔT and ΔQ.

All the results in the following always include dyad and hourly fixed effects (and, when allowed by the model, user fixed effects). Table 21 repeats the analysis in Table 16 by imposing $\Delta T = 30s$ and $\Delta Q = 1\%$ in columns (1-2-3), and $\Delta T = 600s$ and $\Delta Q = 20\%$ in columns (4-5-6). The dependent variable is with fees (1-4), without fees (2-5), and with expected fees (3-6) — further details on this are discussed in Appendix A.2. The results hold whatever the specification. Interestingly, the R^2 is higher for smaller ΔT and ΔQ; larger boundaries could lead to the inclusion of a higher amount of false positives. Table 22 shows that similar considerations apply also when considering alternative specifications of the expertise.

Interesting findings are reported in Table 23, that replicates Table 18 for $\Delta T = 30s$ and $\Delta Q = 1\%$ in columns (1) to (6), and $\Delta T = 600s$ and $\Delta Q = 20\%$ in (7) to (12). The results strengthen the intuition discussed above that larger boundaries reduce the statistical power of the model. Indeed, though the overall pattern is confirmed, the results for $\Delta T = 600s$ and $\Delta Q = 20\%$ are often not statistically significant and the R^2 is smaller. Noteworthy, instead, columns (1) to (6) provide even better results with respect to those reported in the main analysis (with $\Delta T = 300s$ and $\Delta Q = 10\%$), both considering the β_1 coefficients, and the R^2.

These results are especially important, as they demonstrate that the findings described in the main body of the paper are not circumscribed to a specific parametrization of ΔT and ΔQ. The boundaries $\Delta T = 300s$ and $\Delta Q = 10\%$ appear then to be reasonable, though their selection is not backed by a rigorous method but rather by examination of the descriptive statistics. While it seems plausible that smaller intervals lead to higher statistical precision at the cost of excluding some true positives from the sample, larger values of ΔT and ΔQ lead to the inclusion of additional points in the sample, at the cost of a weaker statistical power and likely of a larger fraction of false positives. In addition, I point out that one could incur in self-selection bias by further restricting the boundaries on time and volume. Indeed, as shown in Section 3.3, less skilled users tend to execute actions with larger δT and δQ. I thus prefer an intermediate, more conservative approach, and focus on $\Delta T = 300s$ and $\Delta Q = 10\%$ in the main analysis.

Table 21: Relationship between trade ability and profits. Robustness check with alternative ΔT, ΔQ

Thresholds:		30s, 1%			600s, 20%	
Dep. var.:	With	No	Exp.	With	No	Exp.
	(1)	(2)	(3)	(4)	(5)	(6)
D(Currencies)	1.2287**	0.8634**	1.1033**	1.2304***	0.8557***	1.1568***
	(0.4830)	(0.3784)	(0.4219)	(0.1349)	(0.1221)	(0.1376)
Equiv. $	0.6671	0.3140	0.4420	-0.3636	-1.0275	-0.5332
	(0.8899)	(0.7671)	(0.7590)	(0.8931)	(0.8572)	(0.8843)
Constant	-0.5669	0.2005	-0.6963*	-0.6018***	0.3510***	-0.6794***
	(0.4602)	(0.3622)	(0.4028)	(0.1274)	(0.1153)	(0.1270)
Time FE	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y
N	3935	3935	3935	6554	6554	6554
R-squared	0.71	0.76	0.75	0.64	0.66	0.65

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 22: Relationship between trade ability and profits, alternative proxies. Robustness check with alternative ΔT, ΔQ

Dep. var.:	Spread (with fees)											
Thresholds:	30s, 1%					600s, 20%						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
D(Currencies)	1.229**						1.230***					
	(0.483)						(0.135)					
Log(Currencies)		1.637*						0.634*				
		(0.954)						(0.327)				
Log(Actions)			0.453***						0.270***			
			(0.103)						(0.033)			
D(Metaorder)				0.874***						0.842***		
				(0.241)						(0.146)		
D(Aggressive)					-0.974***						-1.493***	
					(0.344)						(0.173)	
PC1						0.472**						0.188***
						(0.211)						(0.017)
Equiv. $	0.667	1.032	0.493	0.840	0.865	0.845	-0.364	-0.517	1.182	1.133	0.379	1.450
	(0.890)	(0.854)	(0.441)	(0.868)	(0.858)	(0.716)	(0.893)	(1.166)	(1.052)	(1.402)	(1.027)	(1.231)
Constant	-0.567	-2.333	-2.379***	-0.213	0.628***	-2.084*	-0.602***	-0.520	-1.193***	-0.144	0.595***	-0.867***
	(0.460)	(1.735)	(0.690)	(0.243)	(0.054)	(1.226)	(0.127)	(0.566)	(0.223)	(0.122)	(0.040)	(0.142)
Time FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	3935	3935	3935	3935	3935	3935	6554	6554	6554	6554	6554	6554
R-squared	0.71	0.74	0.78	0.71	0.71	0.75	0.64	0.63	0.66	0.63	0.64	0.66

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 23: Responsiveness to official rate variations. Robustness check with alternative ΔT, ΔQ.

Thresholds:	30s, 1%						600s, 20%					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Int. D(Currencies)	1.420						-0.094					
	(1.108)						(3.143)					
Int. Log(Currencies)		3.049*						0.275				
		(1.546)						(0.980)				
Int. Log(Actions)			0.593***						0.040			
			(0.173)						(0.356)			
Int. D(Metaorder)				2.404*						2.024**		
				(1.429)						(0.814)		
Int. D(Aggressive)					-0.608						1.823	
					(1.324)						(2.420)	
Int. PC1						0.704***						0.317
						(0.245)						(0.204)
DeltaER	0.000	-4.226	-2.547	-0.693	1.433	-2.561*	0.390	-0.188	0.035	-1.394	0.246	-2.160
	(.)	(3.074)	(1.524)	(1.024)	(1.121)	(1.430)	(3.226)	(1.782)	(2.410)	(0.965)	(0.667)	(1.694)
Equiv. $	0.012	0.002	-0.021	0.009	0.011	-0.008	-0.935	-0.956	-0.939	-1.037	-0.927	-1.009
	(0.583)	(0.575)	(0.586)	(0.584)	(0.583)	(0.586)	(0.838)	(0.825)	(0.825)	(0.881)	(0.833)	(0.861)
User FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	3923	3923	3923	3923	3923	3923	6322	6322	6322	6322	6322	6322
R-squared	0.81	0.81	0.81	0.81	0.81	0.81	0.71	0.71	0.71	0.71	0.71	0.71

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

A.2 Robustness checks on alternative measures of transaction costs

In the following we repeat the main estimations in Tables 16, 17, 18 for the case in which profits are measured without including the fees, and with the fees a user would expect to pay given his transaction history (Tables 24, 25, 26). Results are consistent with those reported for the main estimations.

Table 24: Relationship between trade ability and profits. Robustness check with alternative measures of profits (no fees, expected fees)

Dep. var.:	Spread (Without fees)				Spread (Expected fees)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
D(Currencies)	1.0687***	1.1651***	0.8727***	0.9316***	1.3977***	1.4436***	1.1583***	1.2160***
	(0.2204)	(0.1880)	(0.1416)	(0.1486)	(0.1845)	(0.1878)	(0.1643)	(0.1702)
Equiv. $	2.3429	1.3837	0.0384	-0.4376	3.7186**	2.6949*	0.4920	0.0296
	(1.6352)	(1.3343)	(1.1974)	(1.1337)	(1.8186)	(1.6110)	(1.2224)	(1.0973)
Constant	0.0823	0.0048	0.2756**	0.2283	-0.9895***	-1.0211***	-0.7191***	-0.7642***
	(0.1582)	(0.1640)	(0.1350)	(0.1428)	(0.1582)	(0.1676)	(0.1540)	(0.1616)
Time FE	N	N	Y	Y	N	N	Y	Y
Dyad FE	N	Y	N	Y	N	Y	N	Y
N	6594	6582	5307	5284	6594	6582	5307	5284
R-squared	0.07	0.12	0.70	0.71	0.12	0.17	0.70	0.70

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 25: Relationship between trade ability and profits, alternative proxies. Robustness check with alternative measures of profits (no fees, expected fees)

Dep. var.:	Spread (Without fees)						Spread (Expected fees)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
D(Currencies)	0.932***						1.216***					
	(0.149)						(0.170)					
Log(Currencies)		0.690**						0.741**				
		(0.289)						(0.338)				
Log(Actions)			0.218***						0.266***			
			(0.044)						(0.040)			
D(Metaorder)				0.101						0.178		
				(0.080)						(0.161)		
D(Aggressive)					-1.119***						-1.449***	
					(0.186)						(0.192)	
PC1						0.147***						0.188***
						(0.030)						(0.034)
Equiv. $	-0.438	-0.462	0.253	-0.563	-0.328	-0.041	0.030	-0.034	0.859	-0.117	0.170	0.530
	(1.134)	(1.217)	(1.170)	(1.279)	(1.122)	(1.231)	(1.097)	(1.230)	(1.216)	(1.294)	(1.090)	(1.270)
Constant	0.228	-0.086	-0.288	1.023***	1.139***	0.030	-0.764***	-0.900	-1.321***	0.238*	0.424***	-0.989***
	(0.143)	(0.509)	(0.290)	(0.072)	(0.027)	(0.237)	(0.161)	(0.595)	(0.265)	(0.139)	(0.029)	(0.266)
Time FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	5284	5284	5284	5284	5284	5284	5284	5284	5284	5284	5284	5284
R-squared	0.71	0.71	0.72	0.70	0.71	0.71	0.70	0.70	0.72	0.69	0.70	0.71

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 26: Responsiveness to official rate variations. Robustness check with alternative measures of profits (no fees, expected fees)

Dep. var.:	Spread (Without fees)						Spread (Expected fees)					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\Delta R \times$D(Currencies)	1.176						1.266					
	(5.242)						(5.226)					
$\Delta R \times$Log(Currencies)		1.880***						1.767**				
		(0.709)						(0.707)				
$\Delta R \times$Log(Actions)			0.162						0.157			
			(0.408)						(0.406)			
$\Delta R \times$D(Metaorder)				2.463***						2.323***		
				(0.854)						(0.824)		
$\Delta R \times$D(Aggressive)					3.416						3.346	
					(3.466)						(3.404)	
$\Delta R \times$PC1						0.523**						0.496**
						(0.216)						(0.212)
ΔR	-0.496	-2.769*	-0.355	-1.248	0.602	-3.182**	-0.607	-2.583*	-0.343	-1.160	0.582	-2.998*
	(5.212)	(1.533)	(2.422)	(0.846)	(0.964)	(1.555)	(5.193)	(1.548)	(2.391)	(0.839)	(0.956)	(1.535)
Equiv. $	-0.730	-0.872	-0.751	-0.810	-0.738	-0.818	-0.705	-0.839	-0.725	-0.780	-0.713	-0.788
	(0.838)	(0.889)	(0.832)	(0.881)	(0.843)	(0.875)	(0.826)	(0.872)	(0.820)	(0.865)	(0.830)	(0.860)
User FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	5142	5142	5142	5142	5142	5142	5142	5142	5142	5142	5142	5142
R-squared	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75	0.75

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

A.3 Robustness checks on learning-by-doing

Our analysis is based on the assumption that trade ability is an innate characteristic for the investors, i.e. it does not increase significantly over time through trading. In Section 3.4 we address the concern that users might instead learn by trading and we show that this is unlikely in our context. As a robustness check, we repeat here the regressions reported in Tables 16, 17, 18 by excluding the users active in multiple markets if the time passed between their first arbitrage action and the first one in a new currency market is large (i.e., more than 14 days). The results are shown in Tables 27, 28, 29, and their interpretation is in accordance with the one reported for the main estimations.

Table 27: Relationship between trade ability and profits. Robustness check on learning-by-doing

Dep. var.:	Spread (with fees)			
	(1)	(2)	(3)	(4)
D(Currencies)	1.6040***	1.6340***	0.8320***	0.8280***
	(0.2099)	(0.2031)	(0.2391)	(0.2850)
Equiv. $	2.4045	1.7055	0.1145	0.5385
	(1.6007)	(1.3292)	(1.1578)	(0.9215)
Constant	-1.0294***	-1.0457***	-0.2440	-0.2314
	(0.1605)	(0.1973)	(0.2202)	(0.2666)
Time FE	N	N	Y	Y
Dyad FE	N	Y	N	Y
N	4817	4805	4053	4032
R-squared	0.21	0.29	0.75	0.76

Standard errors in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 28: Relationship between trade ability and profits, alternative proxies. Robustness check on learning-by-doing

Dep. var.:	Spread (with fees)					
	(1)	(2)	(3)	(4)	(5)	(6)
D(Currencies)	0.8280***					
	(0.2850)					
Log(Currencies)		1.5511**				
		(0.6296)				
Log(Actions)			0.3741***			
			(0.1056)			
D(Metaorder)				1.1473***		
				(0.2191)		
D(Aggressive)					-1.1139***	
					(0.2103)	
PC1						0.2457***
						(0.0832)
Equiv. $	0.5385	0.4517	0.0911	0.5386	0.7176	0.3940
	(0.9215)	(0.8186)	(0.7401)	(0.9046)	(0.8718)	(0.8288)
Constant	-0.2314	-2.3370*	-1.9417***	-0.5109**	0.5793***	-1.3833**
	(0.2666)	(1.1859)	(0.7161)	(0.2200)	(0.0526)	(0.6772)
Time FE	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y
N	4032	4032	4032	4032	4032	4032
R-squared	0.76	0.79	0.81	0.76	0.76	0.78

Standard errors in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

Table 29: Responsiveness to official rate variations. Robustness check on learning-by-doing

Dep. var.:	Spread (with fees)					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta R \times$D(Currencies)	16.616**					
	(6.460)					
$\Delta R \times$Log(Currencies)		3.048				
		(2.129)				
$\Delta R \times$Log(Actions)			1.143**			
			(0.504)			
$\Delta R \times$D(Metaorder)				3.997		
				(4.080)		
$\Delta R \times$D(Aggressive)					-1.333	
					(7.365)	
$\Delta R \times$PC1						0.991
						(0.598)
ΔR	-15.963**	-5.242	-7.116*	-3.198	0.668	-7.463
	(6.042)	(4.754)	(3.605)	(3.883)	(1.580)	(4.947)
Equiv. $	-0.024	-0.051	-0.114	-0.035	-0.035	-0.073
	(0.640)	(0.639)	(0.646)	(0.647)	(0.647)	(0.645)
User FE	Y	Y	Y	Y	Y	Y
Time FE	Y	Y	Y	Y	Y	Y
Dyad FE	Y	Y	Y	Y	Y	Y
N	3957	3957	3957	3957	3957	3957
R-squared	0.84	0.84	0.84	0.84	0.84	0.84

Standard errors in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Errors are clustered at user-level to account for intra-class correlation.

A.4 Analysis of individual fiat-to-fiat markets

Table 30: Profitability of the arbitrage actions, illustrative example on *individual* markets.

			Without fees		With fees	
buy leg	sell leg	official e.r.	implied e.r.	profit/loss	implied e.r.	profit/loss
EUR	USD	1.40	1.405	profit	1.402	profit
EUR	USD	1.40	1.395	loss	1.390	loss
USD	EUR	1.40	1.395	profit	1.399	profit
USD	EUR	1.40	1.405	loss	1.410	loss

Notes: to compare graphically the actions executed in the same currency market, I now construct the profits (spread between implied and official rates) keeping the official rate *fixed*. That is, I compare all actions belonging to a given dyad to the same official rate. Here, for example, I consider the dyad (EUR,USD) and I express implied rates in EURtoUSD, independently of which currency appears in the buy side. When the implied rate is higher than the official one, it is profitable to buy bitcoins in EUR and sell them for USD, that is, to implicitly buy USD using EUR; viceversa when the implied rate is lower. Fees reduce the margins for profit and worsen losses.

Figure 35 focuses on the EUR/USD market; Panel 35a reports the official exchange rate (gray line) along with all the N = 1,931 triangular actions performed in that market, depicted as dots whose color is red when the bitcoins are bought in EUR and sold in USD - that is, actions where users implicitly buy USD using EUR - and vice versa for the green ones. I will refer to them respectively as the 'EUR buy' trades and the 'USD buy' trades. The implied rates in the plots are computed excluding the transaction costs. I underline that the purpose of arbitrage is to *profit* from an asset mispricing across different markets: as shown in Table 30, the points above (below) the line are profitable if the investor buys bitcoins in EUR (USD) and sells for USD (EUR).

Interestingly, I do not observe persistent price deviations with respect to the official exchange rate, nor I observe arbitrage activity taking the form of an organic response to mispricings from multiple individuals independently trading across currency markets in a univocal direction (determined by the sign of the price deviation), pushing prices towards equilibrium. Rather, it is possible to observe that the implied exchange

Figure 35: Comparison between the implied and the official exchange rate for the detected arbitrage actions (left), and the multi-currency trades (right).

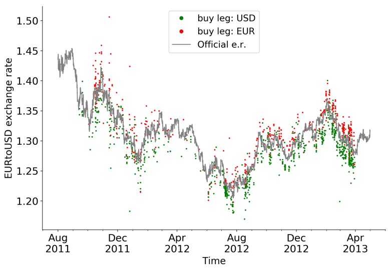

(a) All arbitrage actions

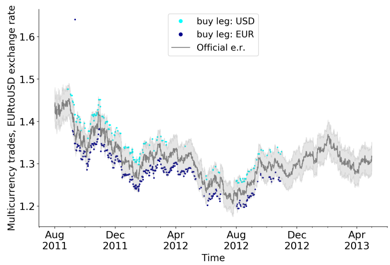

(b) Multi-currency trades, EUR/USD market

Notes: values are reported in absolute terms; the gray line represents the official exchange rate. Note that, while in Panel (a) the trades are distributed coherently with respect to a profit-based logic, the multi-currency trades are always non-profitable and 'anchor' the Bitcoin price across different markets.

rate closely tracks the official one and that actions seemingly exploit opportunities arising due to price volatility between the two of them. I interpret such dynamics as the consequence of many uninformed traders who push prices for a short time far from fundamental values, and the arbitrageurs as the liquidity providers that profit from these demand shocks by exploiting such mispricings. This interpretation is consistent with the results reported in the main body of the paper, with the limited aggressiveness of the expert users on the markets, and with the model of arbitrageurs as liquidity providers described in Gromb and Vayanos (2010).

I hypothesize that the arbitrage opportunities were likely limited by the existence of a feature called multi-currency trading, which was introduced in September 2011 along with the possibility to trade in fiat currencies other than USD. Such feature facilitated the bitcoin trades against minor currencies by allowing the execution of orders involving different currencies in the buy and sell leg, for an additional fee of 2.5%[1]. It is very likely that this mechanism anchored the movements of the bitcoin price in minor markets to the main one, BTCtoUSD. Panel 35b compares the official exchange rate to the implied rate for the trades in EUR and USD that exploited the multi-currency feature. For reasons detailed in Chapter 2 and related to how Mt. Gox internally transcribed the trades, I provide information on this mechanism only with respect to the time period ranging from August 2011 to October 2012 excluding July 2012. These trades are systematically unprofitable. (If I account for and remove the 2.5% fee, the dots converge steadily to the official exchange rate: the gray line shows the official EURtoUSD exchange rate, and the gray band delimits the $\pm 2.5\%$ deviation area.)

Further insights on the dynamics of the arbitrage activity result when I decompose the actions reported in Panel 35a per groups of users. I thus report in Figure 36 the spread between the implied rate and the official rate as a percentage of the latter, and I consider three cases: I show the actions executed by users that made less than 10 arbitrage actions in Panel 36a, and those made by users active in a single market in Panel 36b. In Panel 36c I consider instead only the actions ascribable to the users active in multiple markets. On the right side I include the transaction costs, while on the left side I exclude them; the gray shaded band approximates the area in which the transaction costs might[2] exceed the potential gains originated by the spread between the implied and the official rate. The differences between the first two cases and the third are evident; in the former two, the points are more randomly dispersed around the zero, while in the latter the EUR buy trades and the USD buy trades are clustered. Second, in the latter scenario trades are less dispersed in time, and seemingly part of complex automated strategie, as discussed in Section 3.3.

[1] To conduct such trades, Mt. Gox acted as an intermediary between parties via a virtual internal account. See `https://bit.ly/3jXtTPe`. It is also interesting to notice that only around 1% of the arbitrage actions are composed by legs that are part of multi-currency trades: arbitrageurs do not use this feature, as it entails additional fees.

[2] As discussed in Appendix A.2, the fees paid depend on the users' trading history, thus are specific to each individual and vary across users and in time.

Figure 36: Difference between the implied and the official exchange rate as a percentage of the exchange rate for the EURtoUSD market.

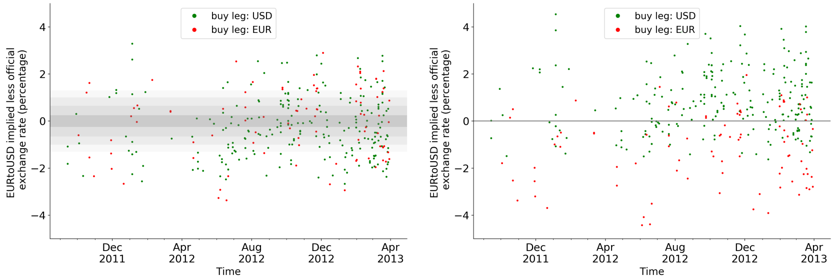

(a) Few actions, without fees (left) & including fees (right)

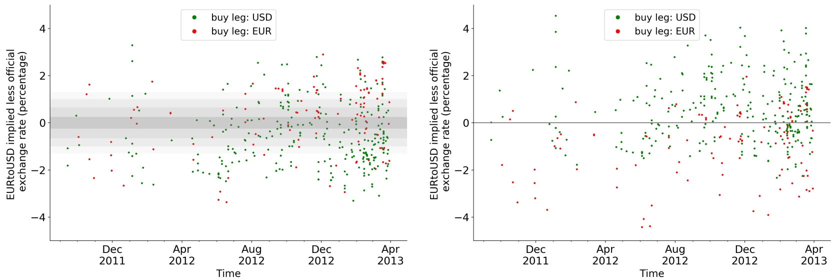

(b) Single market, without fees (left) & including fees (right)

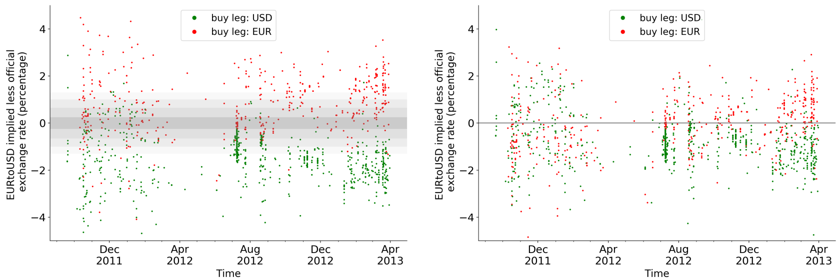

(c) Multiple market, without fees (left) & including fees (right)

Notes: Panel (a) reports only the actions executed by the users that made less than 10 arbitrage actions, Panel (b) those executed by the users active in a single market, and Panel (c) those executed by users active in multiple markets. On the right, fees are included; on the left, fees are excluded. The gray shaded bands represent the area in which transaction costs might exceed the potential profits.

A.5 Supplemental Figures and Tables

Table 31: Summary statistics.

($N = 6,629$)	Mean	St.D.	Min	25%	50%	75%	Max
Profits, fees, %	0.42	1.26	-11.35	0.08	0.62	1.10	18.16
'Equiv. $'	52.54	169.63	0.00	0.36	7.40	41.42	4666.66
ΔR (abs)	0.06	0.08	0.00	0.01	0.04	0.08	1.16
D(Currencies)	0.89	0.31	0.00	1.00	1.00	1.00	1.00
Log(Currencies)	1.64	0.51	0.69	1.10	1.61	1.95	2.48
Log(Actions)	5.94	2.14	0.00	5.30	6.54	7.68	7.68
D(Metaorder)	0.67	0.47	0.00	0.00	1.00	1.00	1.00
D(Aggressive)	0.06	0.23	0.00	0.00	0.00	0.00	1.00
PC1	6.57	2.83	-0.77	5.03	8.16	8.83	8.83

Notes: This table summarizes the main variables employed in the regression analyses. The sample is the number of arbitrage actions (N = 6,629). The profits are expressed as a percentage of the spread between the implied and the official rate, as explained in Table 10; the 'Equiv. $' term is the value of the trade expressed in USD dollars. The remaining rows refer to alternative measures of the investors trading ability. D(Currencies) is equal to 1 if the user exploited multiple currency markets to conduct arbitrage. It is a dummy variable, as well as D(Metaorder) and D(aggressive), respectively equal to 1 if the action was conducted by a user who executed metaorders or aggressive orders. Log(Currencies) is the logarithm of the currency markets exploited by the investor who conducted the arbitrage action, while Log(Actions) is the logarithm of the number of actions they executed. PC1 is the scores of the arbitrage action, obtained by performing a principal component analysis as described in Table 15.

Table 32: Pearson correlation for the main variables used in the model

	$D(Cur)$	$Log(Cur)$	$Log(Act)$	$D(Met)$	$D(Agg)$	PC1	Profits	Eq. $	ΔR
$D(Cur)$	1	0.65	0.8	0.41	-0.44	0.76	0.39	-0.14	-0.10
$Log(Cur)$		1	0.74	0.63	-0.31	0.79	0.22	-0.18	-0.08
$Log(Act)$			1	0.57	-0.53	0.9	0.43	-0.18	-0.11
$D(Met)$				1	-0.35	0.86	0.24	-0.17	-0.04
$D(Agg)$					1	-0.52	-0.28	0.13	0.01
PC1						1	0.39	-0.2	-0.09
Profits							1	-0.01	-0.02
Eq. $								1	0.04
ΔR									1.00

Notes: The correlation is constructed on the sample of the actions (N = 6,629).

Figure 37: Example of an online arbitrage tool (captured by the Internet Archive on 8 April 2013)

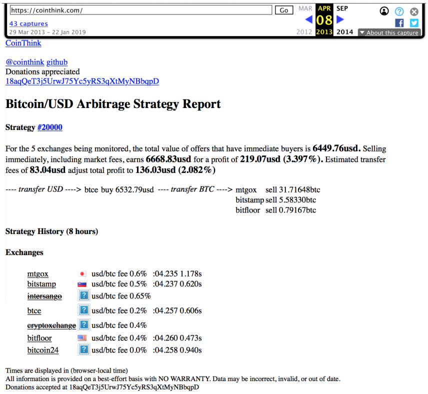

https://cointhink.com/ Go MAR APR SEP
43 captures 08
29 Mar 2013 – 22 Jan 2019 2012 2013 2014 About this capture

CoinThink

@cointhink github
Donations appreciated
18aqQeT3j5UrwJ75Yc5yRS3qXtMyNBbqpD

Bitcoin/USD Arbitrage Strategy Report

Strategy #20000

For the 5 exchanges being monitored, the total value of offers that have immediate buyers is **6449.76usd.** Selling immediately, including market fees, earns **6668.83usd** for a profit of **219.07usd (3.397%).** Estimated transfer fees of **83.04usd** adjust total profit to **136.03usd (2.082%)**

---- *transfer USD* ----> btce buy 6532.79usd ---- *transfer BTC* ----> mtgox sell 31.71648btc
bitstamp sell 5.58330btc
bitfloor sell 0.79167btc

Strategy History (8 hours)

Exchanges

mtgox	usd/btc fee 0.6%	:04.235	1.178s
bitstamp	usd/btc fee 0.5%	:04.237	0.620s
~~intersango~~	usd/btc fee 0.65%		
btce	usd/btc fee 0.2%	:04.257	0.606s
~~cryptoxchange~~	usd/btc fee 0.4%		
bitfloor	usd/btc fee 0.4%	:04.260	0.473s
bitcoin24	usd/btc fee 0.0%	:04.258	0.940s

Times are displayed in (browser-local time)
All information is provided on a best-effort basis with NO WARRANTY. Data may be incorrect, invalid, or out of date.
Donations accepted at 18aqQeT3j5UrwJ75Yc5yRS3qXtMyNBbqpD

Figure 38: Growth of the triangular arbitrage actions as a function of ΔT and ΔQ.

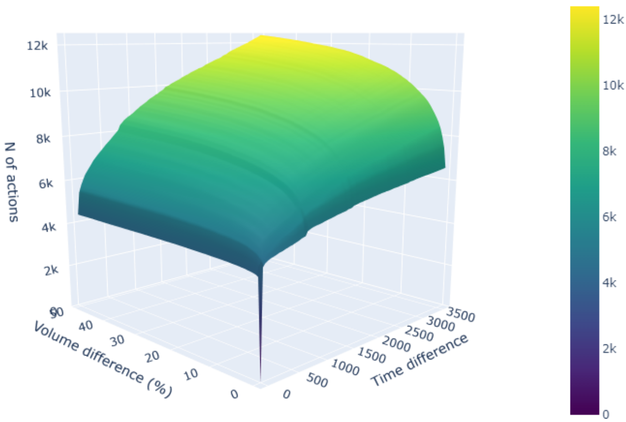

Figure 39: Trading patterns of the 3^{rd} to the 8^{th} user most actives in multiple markets.

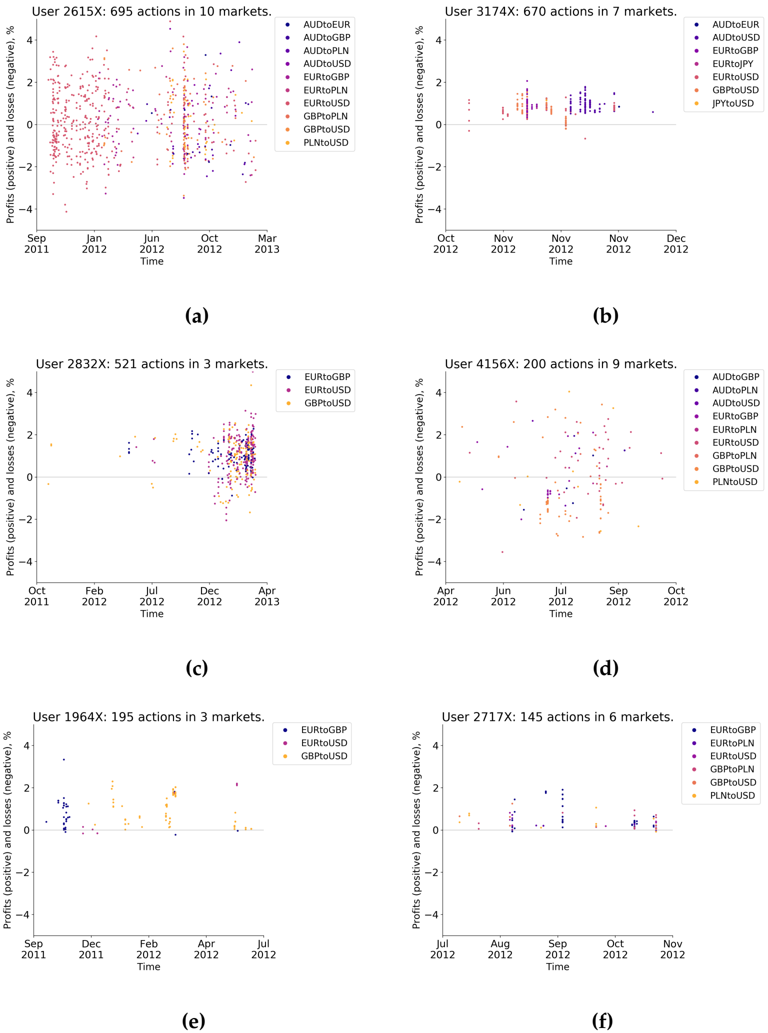

Notes: the y-axis reports the profitability of the actions, depicted as dots, and the x-axis shows their evolution in time. The different colors correspond to actions conducted in different currency markets. I hide the last unit of each user identifier to preserve the anonymity.

Figure 40: Comparison between the implied and the official exchange rate for the detected arbitrage actions (left), and the multi-currency trades (right).

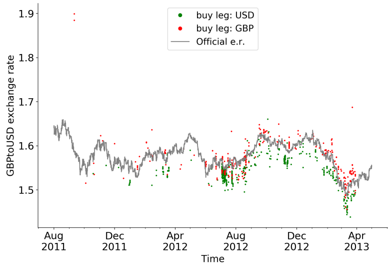

(a) All arbitrage actions, GBPtoUSD market

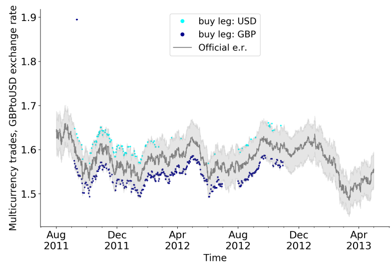

(b) Multi-currency trades, GBPtoUSD market

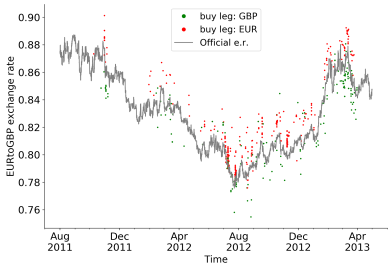

(c) All arbitrage actions, EURtoGBP market

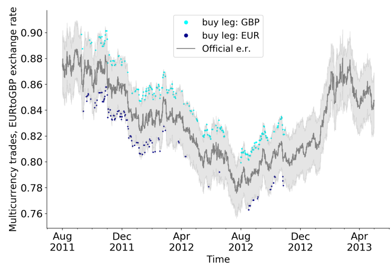

(d) Multi-currency trades, EURtoGBP market

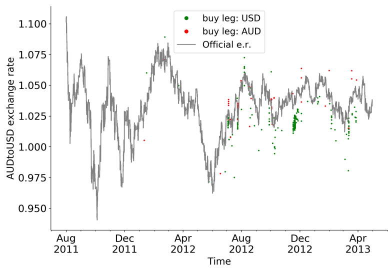

(e) All arbitrage actions, AUDtoUSD market

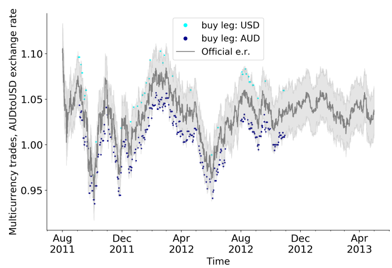

(f) Multi-currency trades, AUDtoUSD market

Notes: the Panels refer respectively to the GBPtoUSD, the EURtoGBP, and the AUDtoUSD markets. Left Panels report the arbitrage actions, right Panels the multi-currency trades.

Figure 41: Difference between the implied and the official exchange rate as a percentage of the exchange rate, actions executed only by users active in multiple markets.

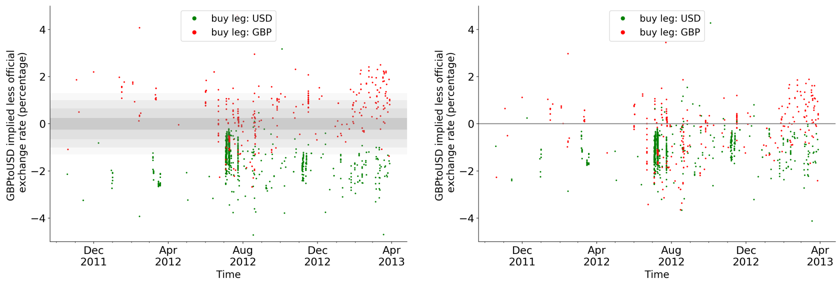

(a) GBPtoUSD. Without fees (left) & including fees (right)

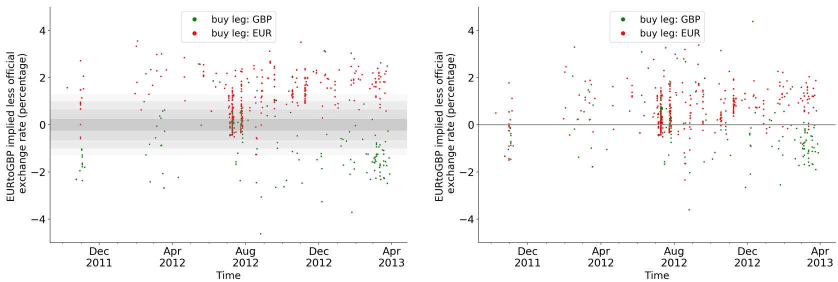

(b) EURtoGBP. Without fees (left) & including fees (right)

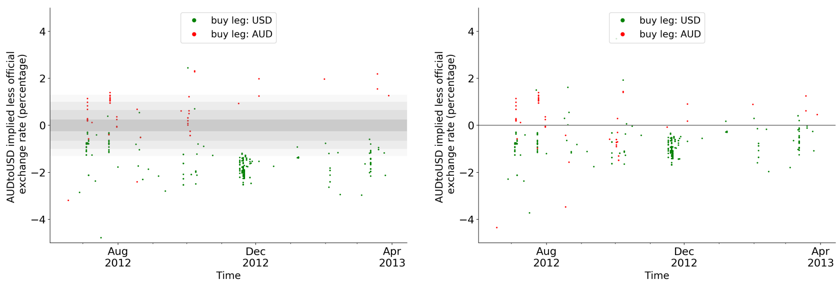

(c) AUDtoUSD. Without fees (left) & including fees (right)

Notes: Panels (a) refer to the GBPtoUSD market, Panels (b) to the EURtoGBP, and Panels (c) to the AUDtoUSD market. I report only the actions of users active in many markets as the other cases are negligible. On the right, fees are included; on the left, fees are excluded.

Appendix B

Supplementary material on two-point arbitrage

B.1 Statistics on the choice of the parameters

This appendix reports statistics related to Sections 4.3.1 and 4.3.2 and additional information useful to better understand the methodology I proposed to identify the arbitrageurs.

Figure 42 shows the 15 days (centered) moving average for the number of daily trades that were executed in Bitstamp, BTC-e, and Mt. Gox (using both the Bitcoincharts.com data and those of the leaked dataset): the users' trading activity is definitely higher in the latter, where the number of daily trades is often bigger by one order of magnitude or more. As a consequence, I hypothesize that users performing arbitrage would necessarily do it exploiting Mt. Gox as the main platform, due to the higher market liquidity.

Second, I address the choice of aggregating the trades at the hourly level. The trades in the counterpart exchanges are sparse across the day: thus, I believe that constructing sequences of actions at time intervals shorter than one hour for the similarity metrics (e.g., 10-minute or 5-minute time windows) would increase the complexity of the approach without adding explanatory information. In addition, Figure 43 sows the evolution in time of the coefficient of variation for the prices in the counterpart exchanges, aggregated at hourly level. Prices do not vary consistently in the time window considered, thus I can safely assume that prices are 'fixed' and do not vary within the same hourly window.

Figure 42: Centered 15 days moving average for the number of daily trades executed in BTC-e, Bitstamp, and Mt. Gox.

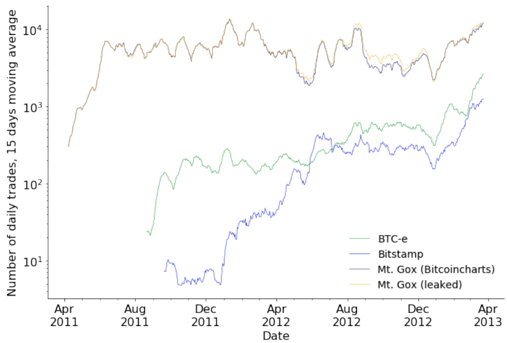

Figure 43: Coefficient of variation for the hourly price in Bitstamp (left) and BTC-e (right).

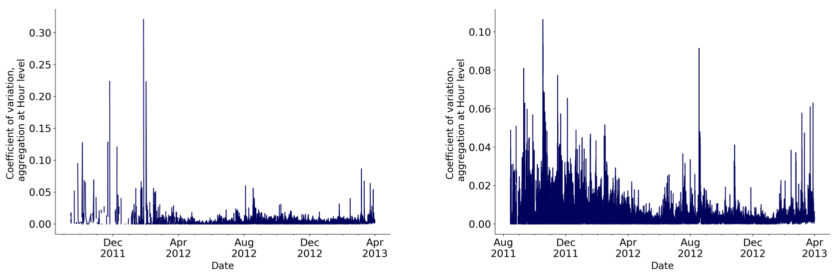

(a) **(b)**

Table 33: Quantiles for the Bitstamp and BTC-e users, relative to the number of hours of activity.

	Hours of activity			Hours of activity	
	Bitstamp	BTC-e		Bitstamp	BTC-e
mean	29.18	26.29	**50%**	5	5
std	170.61	159.51	**75%**	16	14
min	1	1	**95%**	97	86
25%	2	2	**max**	11352	11352

Table 33 reports statistics (mean, standard deviation, quantiles) on the number of hours of activity for the Mt. Gox users. As the data show, the assumption that arbitrageurs are users who performed arbitrage over a long period of time leads to a significant trade-off between the size of the sample analyzed and the accuracy of the results. In Chapter 3 it is shown that the arbitrageurs that executed most of the arbitrage actions are few and large users. For this reason I assume it is safe to restrict the analysis to a limited group of users who executed many trades. Second, as a validation, I repeat the analysis with a larger sample (instead of 100, I consider users active for more than 10 hours) and I report the results in Appendix B.2. I do not increase further the sample size, as for users with too few actions it would be difficult to distinguish whether positive outcomes are due to random fluctuations or to actual strategic actions. In the latter, I accept a higher number of false positives in the first stage, expecting them to be ruled out in the post-filtering phase.

In Section 4.3.1 I introduced a cutoff on the share of arbitrage-like actions, to divide users in two sets (potential arbitrageurs and non arbitrageurs) and to inspect the differences between them. In order to provide a criterion to choose that parameter, I compute the Pearson correlation between the price differences of Bitstamp and BTC-e against Mt. Gox: in principle, since an investor can be identified as a potential arbitrageur in both exchanges, if the ideal series are strongly correlated a higher cutoff is needed to better classify users; since the value is low (0.014), I imposed the cutoff at 0.33 (that is, 33% of arbitrage-like actions);

I repeated the analysis with a cutoff at 0.5. Results are similar; here I included only the results for 0.33. Note anyway that this parameter serves mainly for qualitative purposes: it is not used in the post-filtering procedure to detect the final set of arbitrageurs. In that step of the analysis I use the same threshold (0.3) both for the share of arbitrage-like actions and for the share of mirror actions. I explicitly choose a large threshold in order to maintain a conservative approach in the arbitrageurs identification.

Below I report additional information for the two sets of users, namely the potential arbitrageurs and the non-arbitrageurs, to highlight the differences between their distributions of the shares of actions with a mirror trade per user. I follow the scheme adopted in the main analysis (Bitstamp on the left, BTC-e on the right, restrictive scenario on top and looser scenario in the bottom part). I report the kernel density estimation together with statistics on the mean, the standard deviation, the skewness and the kurtosis (Figure 44), qq-plots (Figure 45) and pp-plots (Figure 46). I plotted the probability plots following the methodology proposed in Gibbons and Chakraborti (2011) and Thode (2002).

Finally, the parameters ΔT and ΔQ are chosen based on the empirical evidence reported in Chapter 3. The vast majority of trades lie in a small interval around $\Delta T = 0s$ and $\Delta Q = 0\%$. As in the previous Chapter I considered two different scenarios, a more conservative and a less conservative one. I follow the same approach and choose the parameters accordingly.

B.2 Analysis for users active for more than 10 hours

In the following paragraphs I repeat, as a robustness check, the analyses reported in Section 4.3 by considering the users that were active for more than 10 hours instead of 100. Figures 47,48, and 49 show that the results are qualitatively similar to those reported in the main analysis. In total, 821 arbitrageurs are identified, 192 between Mt. Gox and Bitstamp, and 759 between Mt. Gox and BTC-e.

Figure 44: Kernel density estimation.

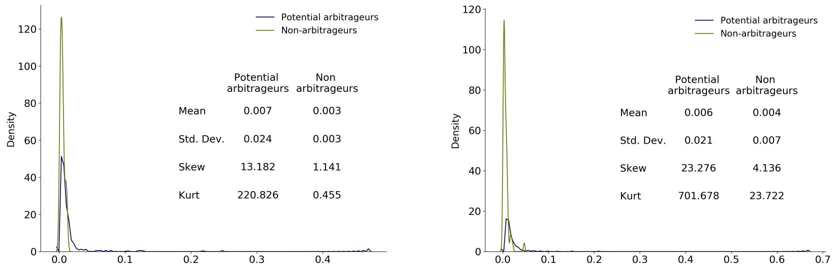

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

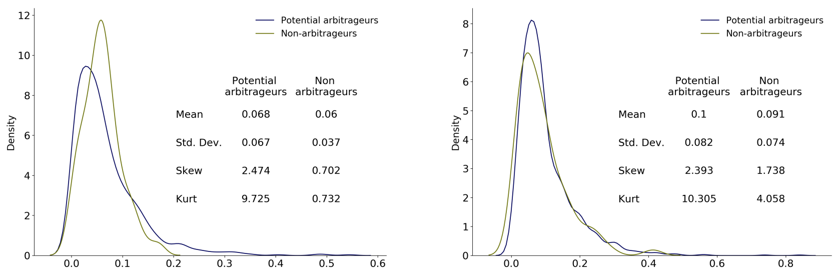

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Figure 45: QQ-plots.

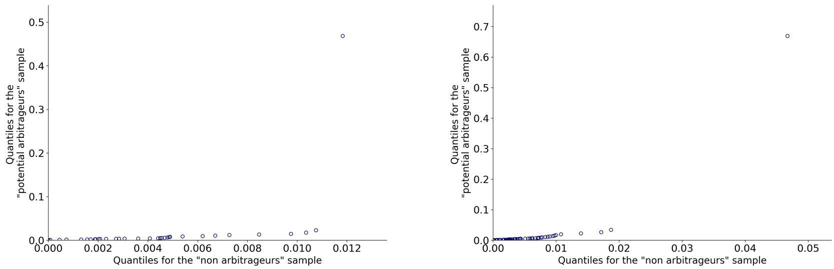

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

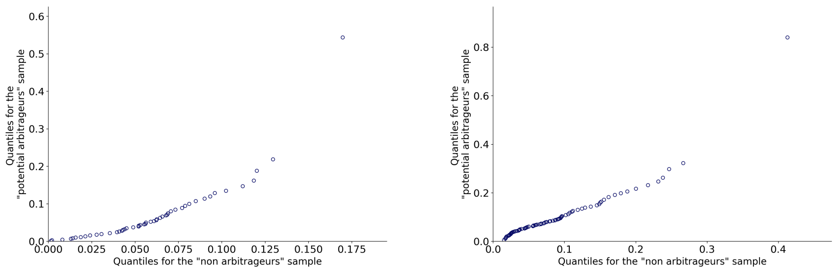

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Figure 46: PP-plots.

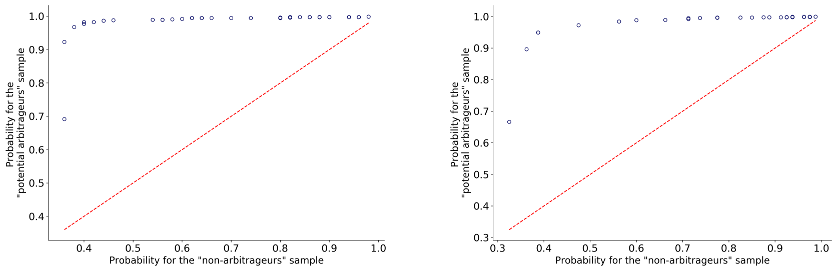

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

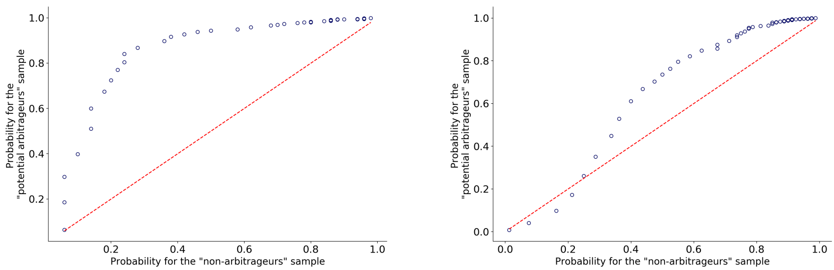

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Figure 47: Share of trades that correspond to two-point arbitrage actions with another exchange, per user, 10 hours or more.

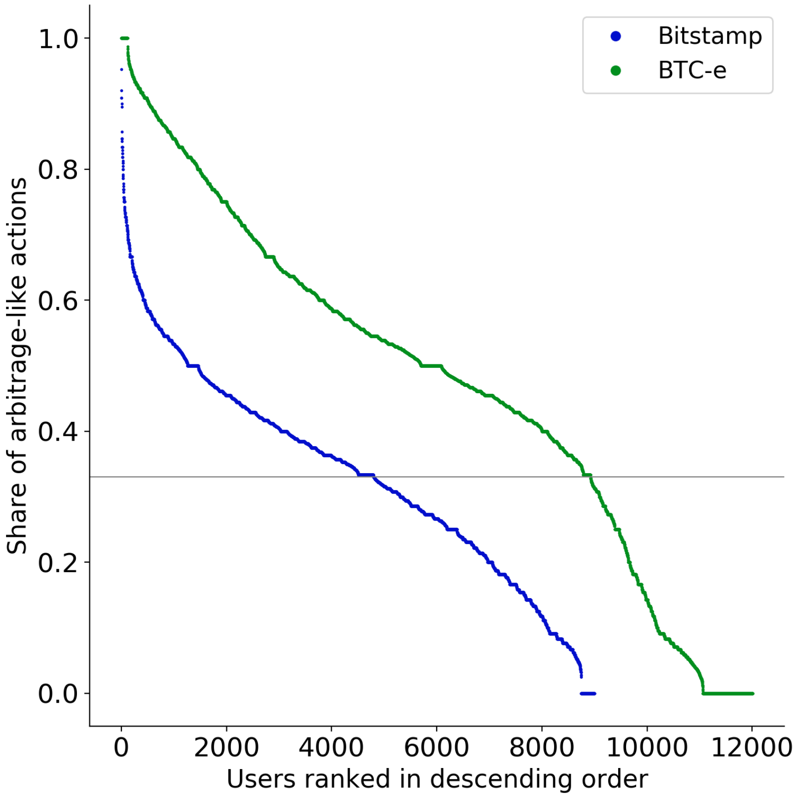

Figure 48: Share of actions with a mirror trade in the counter-part exchange, 10 hours or more.

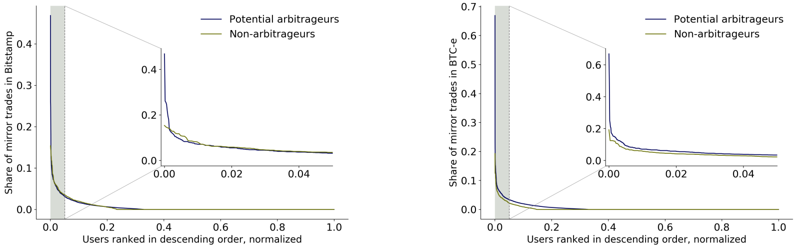

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

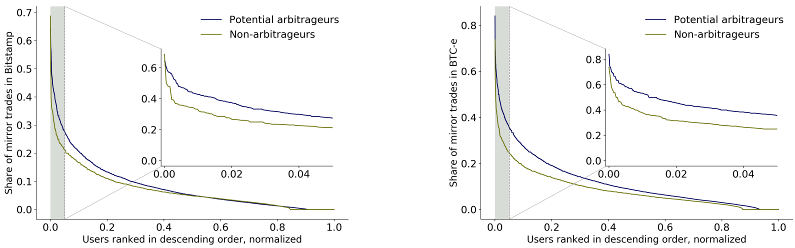

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Notes: users for which the matching occurs on 10 hours or less are excluded from this analysis. Share of actions with a mirror trade in the counterpart exchange (left: Bitstamp, right: BTC-e), per user ranked in descending order. Users are represented on two different lines (yellow: non arbitrageurs, blue: potential arbitrageurs), to qualitatively inspect the differences between the two groups. Panel (a) reports the results for a more restrictive scenario, panel (b) for a less restrictive one.

Figure 49: Post-filtering procedure, 10 hours or more.

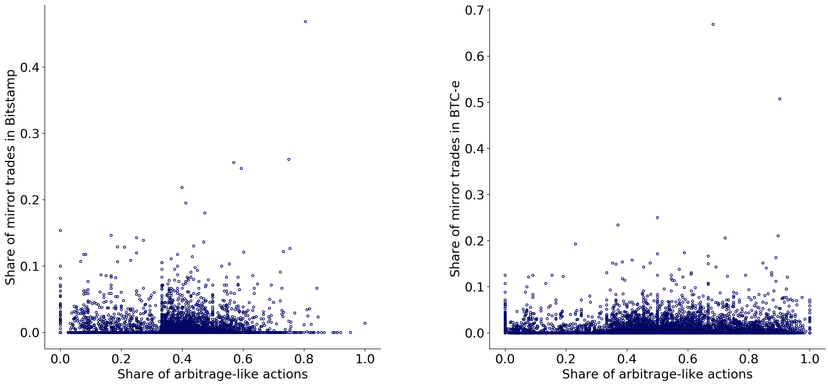

(a) more restrictive scenario: $\Delta T = 30s$ and $\Delta Q = 1\%$

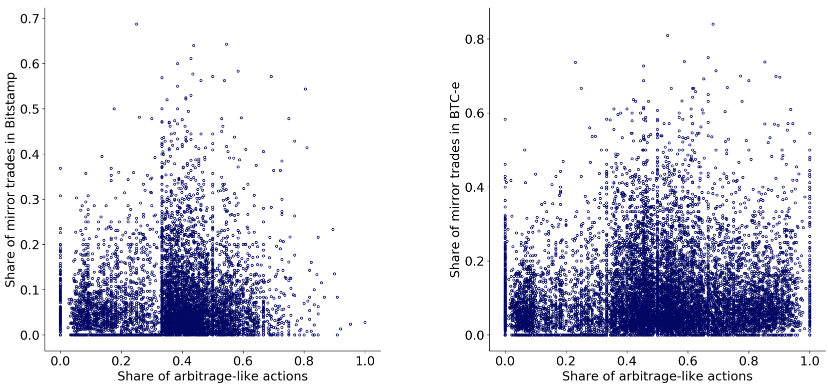

(b) less restrictive scenario: $\Delta T = 300s$ and $\Delta Q = 10\%$

Notes: I plot a dot per each user for which the matching occurs on more than 10 hours. The y-axis is the share of actions with a mirror trade in the counterpart exchange, and the x-axis is the share of arbitrge-like actions (left: Bitstamp, right: BTC-e). Again, Panel (a) reports the results for a more restrictive scenario, and Panel (b) for a looser one.

Bibliography

Aliber, Robert Z (1973). “The Interest Rate Parity Theorem: A Reinterpretation”. In: *Journal of Political Economy* 81.6, pp. 1451–1459.

Allen, Sarah et al. (2020). *Design Choices for Central Bank Digital Currency: Policy and Technical Considerations*. Tech. rep. National Bureau of Economic Research.

Amihud, Yakov and Haim Mendelson (1987). “Trading mechanisms and stock returns: An empirical investigation”. In: *The Journal of Finance* 42.3, pp. 533–553.

— (1991). “Volatility, efficiency, and trading: Evidence from the Japanese stock market”. In: *The Journal of Finance* 46.5, pp. 1765–1789.

Anderson, Ross John et al. (2018). “Bitcoin redux”. In:

Androulaki, Elli et al. (2013). “Evaluating user privacy in bitcoin”. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 34–51.

Antonopoulos, Andreas M (2017). *Mastering Bitcoin: Programming the open blockchain*. " O’Reilly Media, Inc."

Antonopoulos, Andreas M and Gavin Wood (2018). *Mastering ethereum: building smart contracts and dapps*. O’reilly Media.

Athey, Susan et al. (2016). “Bitcoin pricing, adoption, and usage: Theory and evidence”. In:

Back, Adam (2002). *Hashcash - a denial of service countermeasure*.

Badev, Anton I and Matthew Chen (2014). “Bitcoin: Technical background and data analysis”. In:

Barber, Simon et al. (2012). “Bitter to better—how to make bitcoin a better currency”. In: *International conference on financial cryptography and data security*. Springer, pp. 399–414.

Barberis, Nicholas and Richard Thaler (2003). “A survey of behavioral finance”. In: *Handbook of the Economics of Finance* 1, pp. 1053–1128.

Bariviera, Aurelio F (2017). "The inefficiency of Bitcoin revisited: A dynamic approach". In: *Economics Letters* 161, pp. 1–4.
Bech, Morten L and Rodney Garratt (2017). "Central bank cryptocurrencies". In: *BIS Quarterly Review September*.
Bistarelli, Stefano et al. (2019). "Model-based arbitrage in multi-exchange models for Bitcoin price dynamics". In: *Digital Finance* 1.1-4, pp. 23–46.
Bjerg, Ole (2017). "Designing new money-the policy trilemma of central bank digital currency". In:
Björk, Tomas (2009). *Arbitrage theory in continuous time*. Oxford university press.
Black, Fischer (1986). "Noise". In: *The journal of finance* 41.3, pp. 528–543.
Black, Fischer and Myron Scholes (1973). "The Pricing of Options and Corporate Liabilities". In: *Journal of Political Economy* 81, pp. 637–654.
Bodie, Zvi, Alex Kane, and Alan Marcus (2018). *Investments*. McGraw-Hill Education.
Böhme, Rainer et al. (2015). "Bitcoin: Economics, Technology, and Governance". In: *Journal of Economic Perspectives* 29.2, pp. 213–238.
Bonneau, Joseph et al. (2015). "Sok: Research perspectives and challenges for bitcoin and cryptocurrencies". In: *2015 IEEE symposium on security and privacy*. IEEE, pp. 104–121.
Bouri, Elie, Georges Azzi, and Anne H Dyhrberg (2016). *On the return–volatility relationship in the Bitcoin market around the price crash of 2013*. Tech. rep. Economics Discussion Papers, No. 2016–41, Kiel Institute for the World Economy (IfW), Kiel.
Brandvold, Morten et al. (2015). "Price discovery on Bitcoin exchanges". In: *Journal of International Financial Markets, Institutions and Money* 36, pp. 18–35.
Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan (2014). "High-frequency trading and price discovery". In: *The Review of Financial Studies* 27.8, pp. 2267–2306.
Budish, Eric, Peter Cramton, and John Shim (2015). "The high-frequency trading arms race: Frequent batch auctions as a market design response". In: *The Quarterly Journal of Economics* 130.4, pp. 1547–1621.
Buterin, Vitalik et al. (2014). *A next-generation smart contract and decentralized application platform*. Whitepaper published online at `https://ethereum.org/en/whitepaper/`. Accessed on 29 October 2020.
Buterin, Vitalik (2017). *The meaning of decentralization*. Online at Medium.com. `https://bit.ly/3mBu83B`. Accessed on 29 October 2020.

Catalini, Christian and Joshua S Gans (2016). *Some simple economics of the blockchain*. Tech. rep. National Bureau of Economic Research.
Cerny, Ales (2009). *Mathematical techniques in finance: tools for incomplete markets*. Princeton University Press.
Ceruleo, Paola (2014). "Bitcoin: a rival to fiat money or a speculative financial asset?" In:
Chaum, David (1983). "Blind signatures for untraceable payments". In: *Advances in cryptology*. Springer, pp. 199–203.
Chaum, David, Amos Fiat, and Moni Naor (1988). "Untraceable electronic cash". In: *Conference on the Theory and Application of Cryptography*. Springer, pp. 319–327.
Cheah, Eng-Tuck et al. (2018). "Long memory interdependency and inefficiency in Bitcoin markets". In: *Economics Letters* 167, pp. 18–25.
Chen, Weili et al. (2019). "Market manipulation of bitcoin: Evidence from mining the Mt. Gox transaction network". In: *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, pp. 964–972.
Cheung, Adrian, Eduardo Roca, and Jen-Je Su (2015). "Crypto-currency Bubbles: An Application of the Phillips–Shi–Yu (2013) Methodology on Mt. Gox Bitcoin Prices". In: *Applied Economics* 47.23, pp. 2348–2358.
Christie, William G and Paul H Schultz (1994). "Why do NASDAQ market makers avoid odd-eighth quotes?" In: *The Journal of Finance* 49.5, pp. 1813–1840.
Christin, Nicolas (2013). "Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace". In: *Proceedings of the 22nd international conference on World Wide Web*, pp. 213–224.
Ciaian, Pavel, Miroslava Rajcaniova, and d'Artis Kancs (2016). "The economics of BitCoin price formation". In: *Applied Economics* 48.19, pp. 1799–1815.
Clark-Murphy, Marilyn and Geoffrey Soutar (2005). "Individual investor preferences: A segmentation analysis". In: *The Journal of Behavioral Finance* 6.1, pp. 6–14.
Conti, Mauro et al. (2018). "A survey on security and privacy issues of bitcoin". In: *IEEE Communications Surveys & Tutorials* 20.4, pp. 3416–3452.
Cox, John C and Stephen A Ross (1976). "The valuation of options for alternative stochastic processes". In: *Journal of financial economics* 3.1-2, pp. 145–166.
Cox, John C, Stephen A Ross, and Mark Rubinstein (1979). "Option pricing: A simplified approach". In: *Journal of financial Economics* 7.3, pp. 229–263.

CPMI (2015). *Digital currencies*. Tech. rep. Committee on Payments and Market Infrastructures, BIS.

Crépellière, Tommy and Stefan Zeisberger (2020). "Arbitrage in the Market for Cryptocurrencies". In: *Available at SSRN 3606053*.

Dai, Wei (1998). *B-money*. Website-Blog. See https://nakamotoinstitute.org/b-money/.

Daian, Philip et al. (Apr. 2019). "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges". http://arxiv.org/abs/1904.05234 (last access: 7 May 2019).

De Bortoli, Daiane et al. (2019). "Personality traits and investor profile analysis: A behavioral finance study". In: *PloS one* 14.3, e0214062.

De Filippi, Primavera (2014). "Bitcoin: a regulatory nightmare to a libertarian dream". In: *Internet Policy Review* 3.2.

— (2018). *Blockchain and the law: The rule of code*. Harvard University Press.

De Jong, Frank and Barbara Rindi (2009). *The microstructure of financial markets*. Cambridge University Press.

De Long, J Bradford et al. (1990). "Noise trader risk in financial markets". In: *Journal of political Economy* 98.4, pp. 703–738.

Decker, Christian and Roger Wattenhofer (2013). "Information propagation in the bitcoin network". In: *IEEE P2P 2013 Proceedings*. IEEE, pp. 1–10.

Delbaen, Freddy and Walter Schachermayer (1994). "A general version of the fundamental theorem of asset pricing". In: *Mathematische annalen* 300.1, pp. 463–520.

— (1998). "The fundamental theorem of asset pricing for unbounded stochastic processes". In: *Mathematische annalen* 312, pp. 215–250.

— (2006). *The mathematics of arbitrage*. Springer Science & Business Media.

Diffie, Whitfield and Martin Hellman (1976). "New directions in cryptography". In: *IEEE transactions on Information Theory* 22.6, pp. 644–654.

Dimitri, Nicola (2019). "Transaction Fees, Block Size Limit, and Auctions in Bitcoin". In: *Ledger* 4.

Dong, Huijian and Weiguo Dong (2015). "Bitcoin: Exchange rate parity, risk premium, and arbitrage stickiness". In:

Donier, Jonathan and Julius Bonart (2015). "A million metaorder analysis of market impact on the Bitcoin". In: *Market Microstructure and Liquidity* 1.02, p. 1550008.

Duffie, Darrell (2010). *Dynamic asset pricing theory*. Princeton University Press.
Dwork, Cynthia and Moni Naor (1992). “Pricing via processing or combatting junk mail”. In: *Annual International Cryptology Conference*. Springer, pp. 139–147.
Dybvig, Philip H and Stephen Ross (1989). “Arbitrage”. In: *Finance*. Springer, pp. 57–71.
Dyhrberg, Anne H, Sean Foley, and Jiri Svec (2018). “How investible is Bitcoin? Analyzing the liquidity and transaction costs of Bitcoin markets”. In: *Economics Letters* 171, pp. 140–143.
Easley, David, Maureen O’Hara, and Soumya Basu (2019). “From mining to markets: The evolution of bitcoin transaction fees”. In: *Journal of Financial Economics* 134.1, pp. 91–109.
Eyal, Ittay and Emin Gün Sirer (2014). “Majority is not enough: Bitcoin mining is vulnerable”. In: *International conference on financial cryptography and data security*. Springer, pp. 436–454.
Fama, Eugene (1970). “Efficient Capital Markets: A Review of Theory and Empirical Work”. In: *Journal of Finance* 25.2, pp. 383–417.
Feder, Amir, Neil Gandal, JT Hamrick, and Tyler Moore (2018). “The impact of DDoS and other security shocks on Bitcoin currency exchanges: Evidence from Mt. Gox”. In: *Journal of Cybersecurity* 3.2, pp. 137–144.
Feder, Amir, Neil Gandal, JT Hamrick, Tyler Moore, et al. (2018). *The Economics of Cryptocurrency Pump and Dump Schemes*. Tech. rep. CEPR Discussion Papers.
Fischer, Thomas Günter, Christopher Krauss, and Alexander Deinert (2019). “Statistical arbitrage in cryptocurrency markets”. In: *Journal of Risk and Financial Management* 12.1, p. 31.
Foley, Sean, Jonathan R Karlsen, and Tālis J Putniņš (2019). “Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?” In: *The Review of Financial Studies* 32.5, pp. 1798–1853.
Foucault, Thierry (1999). “Order flow composition and trading costs in a dynamic limit order market”. In: *Journal of Financial markets* 2.2, pp. 99–134.
Froot, Kenneth A and Emil M Dabora (1999). “How are stock prices affected by the location of trade?” In: *Journal of financial economics* 53.2, pp. 189–216.
Gandal, Neil and Hanna Halaburda (2014). “Competition in the cryptocurrency market”. In:

Gandal, Neil, JT Hamrick, et al. (2018). "Price manipulation in the Bitcoin ecosystem". In: *Journal of Monetary Economics* 95, pp. 86–96.
Garay, Juan, Aggelos Kiayias, and Nikos Leonardos (2015). "The bitcoin backbone protocol: Analysis and applications". In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, pp. 281–310.
Garcia, David et al. (2014). "The digital traces of bubbles: Feedback cycles between socio-economic signals in the Bitcoin economy". In: *Journal of the Royal Society Interface* 11.99.
Gibbons, Jean Dickinson and Subhabrata Chakraborti (2011). *Nonparametric statistical inference*. Springer.
Glaser, Florian et al. (2014). "Bitcoin-asset or currency? Revealing users' hidden intentions". In:
Glosten, Lawrence R (1994). "Is the electronic open limit order book inevitable?" In: *The Journal of Finance* 49.4, pp. 1127–1161.
Gramoli, Vincent (2020). "From blockchain consensus back to byzantine consensus". In: *Future Generation Computer Systems* 107, pp. 760–769.
Griffin, John M and Amin Shams (2020). "Is Bitcoin really untethered?" In: *The Journal of Finance* 75.4, pp. 1913–1964.
Gromb, Denis and Dimitri Vayanos (2010). "Limits of arbitrage". In: *Annu. Rev. Financ. Econ.* 2.1, pp. 251–275.
Grossman, Sanford J and Joseph E Stiglitz (1980). "On the impossibility of informationally efficient markets". In: *The American economic review* 70.3, pp. 393–408.
Haber, Stuart and W Scott Stornetta (1990). "How to time-stamp a digital document". In: *Conference on the Theory and Application of Cryptography*. Springer, pp. 437–455.
— (1997). "Secure names for bit-strings". In: *Proceedings of the 4th ACM Conference on Computer and Communications Security*, pp. 28–35.
Halaburda, Hanna, Miklos Sarvary, et al. (2016). "Beyond bitcoin". In: *The Economics of Digital Currencies*.
Harris, Larry (2003). *Trading and exchanges: Market microstructure for practitioners*. OUP USA.
Harris, Lawrence and Eitan Gurel (1986). "Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures". In: *the Journal of Finance* 41.4, pp. 815–829.
Harrison, J. Michael and David M. Kreps (1979). "Martingales and Arbitrage in Multiperiod Security Markets". In: *Journal of Economic Theory* 20, pp. 381–408.

Harrison, J. Michael and Stanley R. Pliska (1981). "Martingales and stochastic integrals in the theory of continuous trading". In: *Stochastic processes and their applications* 11.3, pp. 215–260.

Haryanto, Steven, Athor Subroto, and Maria Ulpah (2019). "Disposition effect and herding behavior in the cryptocurrency market". In: *Journal of Industrial and Business Economics*, pp. 1–18.

Hasbrouck, Joel and Gideon Saar (2013). "Low-latency trading". In: *Journal of Financial Markets* 16.4, pp. 646–679.

Hattori, Takahiro and Ryo Ishida (2018). "Do Investors Arbitrage in the Cryptocurrency Market? Evidence from the Bitcoin Futures Market". In:

Hautsch, Nikolaus, Christoph Scheuch, and Stefan Voigt (2018). "Limits to arbitrage in markets with stochastic settlement latency". In:

Hendershott, Terrence, Charles M Jones, and Albert J Menkveld (2011). "Does algorithmic trading improve liquidity?" In: *The Journal of finance* 66.1, pp. 1–33.

Hirano, Yuki et al. (2018). "Analysis Of Bitcoin Market Efficiency By Using Machine Learning". In: *CBU International Conference Proceedings*. Vol. 6, pp. 175–180.

Hougan, Matthew et al. (2019). *Economic and Non-Economic Trading in Bitcoin: Exploring the Real Spot Market For The World's First Digital Commodity*. Tech. rep. Bitwise Asset Management.

Ingersoll, Jonathan E (1987). *Theory of financial decision making*. Vol. 3. Rowman & Littlefield.

Johnson, Barry (2010). *Algorithmic Trading & DMA: An introduction to direct access trading strategies*. Vol. 1. 4Myeloma Press London.

Johnson, Benjamin et al. (2014). "Game-theoretic analysis of DDoS attacks against Bitcoin mining pools". In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 72–86.

Jones, Charles M, Gautam Kaul, and Marc L Lipson (1994). "Transactions, volume, and volatility". In: *The Review of Financial Studies* 7.4, pp. 631–651.

Katsiampa, Paraskevi (2017). "Volatility estimation for Bitcoin: A comparison of GARCH models". In: *Economics Letters* 158, pp. 3–6.

Kim, Thomas (2017). "On the transaction cost of Bitcoin". In: *Finance Research Letters* 23, pp. 300–305.

Kirilenko, Andrei et al. (2017). "The flash crash: High-frequency trading in an electronic market". In: *The Journal of Finance* 72.3, pp. 967–998.

Kirilenko, Andrei A and Gui Lamacie (2015). "Latency and asset prices". In: *Available at SSRN 2546567*.

Kondor, Peter (2009). “Risk in dynamic arbitrage: the price effects of convergence trading”. In: *The Journal of Finance* 64.2, pp. 631–655.

Kourtidis, Dimitrios, Željko Šević, and Prodromos Chatzoglou (2011). “Investors’ trading activity: A behavioural perspective and empirical results”. In: *The Journal of Socio-Economics* 40.5, pp. 548–557.

Kreps, David M. (1981). “Arbitrage and equilibrium in economies with infinitely many commodities”. In: *Journal of Mathematical Economics* 8.1, pp. 15–35.

Kristoufek, Ladislav (2018). “On Bitcoin markets (in) efficiency and its evolution”. In: *Physica A: Statistical Mechanics and its Applications* 503, pp. 257–262.

Kroeger, Alexander and Asani Sarkar (2017). “Law of one bitcoin price?” In: *Federal Reserve Bank of Philadelphia*.

Kroll, Joshua A, Ian C Davey, and Edward W Felten (2013). “The economics of Bitcoin mining, or Bitcoin in the presence of adversaries”. In: *Proceedings of WEIS*. Vol. 2013, p. 11.

Krückeberg, Sinan and Peter Scholz (2020). “Decentralized Efficiency? Arbitrage in bitcoin Markets”. In: *Financial Analysts Journal*.

Lamont, Owen A and Richard H Thaler (2003a). “Anomalies: The law of one price in financial markets”. In: *Journal of Economic Perspectives* 17.4, pp. 191–202.

— (2003b). “Can the market add and subtract? Mispricing in tech stock carve-outs”. In: *Journal of Political Economy* 111.2, pp. 227–268.

Lamport, Leslie, Robert Shostak, and Marshall Pease (1982). “The Byzantine Generals Problem”. In: *ACM Transactions on Programming Languages and Systems* 4.3, pp. 382–401.

Lee, Charles MC and Balkrishna Radhakrishna (2000). “Inferring investor behavior: Evidence from TORQ data”. In: *Journal of Financial Markets* 3.2, pp. 83–111.

Lee, Charles MC and Mark J Ready (1991). “Inferring trade direction from intraday data”. In: *The Journal of Finance* 46.2, pp. 733–746.

Lee, Seungho, Nabil El Meslmani, and Lorne N Switzer (2020). “Pricing Efficiency and Arbitrage in the Bitcoin Spot and Futures Markets”. In: *Research in International Business and Finance* 53.

Leung, Tim and Hung Nguyen (2019). “Constructing cointegrated cryptocurrency portfolios for statistical arbitrage”. In: *Studies in Economics and Finance*.

Lintner, John (1965). “Security Prices, Risk and Maximal Gains from Diversification”. In: *Journal of Finance* 20.4, pp. 587–615.

Lischke, Matthias and Benjamin Fabian (2016). "Analyzing the bitcoin network: The first four years". In: *Future Internet* 8.1, p. 7.
Makarov, Igor and Antoinette Schoar (2020). "Trading and arbitrage in cryptocurrency markets". In: *Journal of Financial Economics* 135.2, pp. 293–319.
Malkiel, Burton G (2003). "The efficient market hypothesis and its critics". In: *Journal of economic perspectives* 17.1, pp. 59–82.
Meiklejohn, Sarah et al. (2013). "A fistful of bitcoins: characterizing payments among men with no names". In: *Proceedings of the 2013 conference on Internet measurement conference*, pp. 127–140.
Merkle, Ralph C (1980). "Protocols for public key cryptosystems". In: *1980 IEEE Symposium on Security and Privacy*. IEEE, pp. 122–122.
— (1987). "A digital signature based on a conventional encryption function". In: *Conference on the theory and application of cryptographic techniques*. Springer, pp. 369–378.
Merton, Robert (1987). "A Simple Model of Capital Market Equilibrium with Incomplete Information". In: *Journal of Finance* 42.3, pp. 483–510.
Merton, Robert C (1973). "Theory of rational option pricing". In: *The Bell Journal of economics and management science*, pp. 141–183.
Mishkin, Frederic S (2018). *The economics of money, banking, and financial markets, global edition*. Pearson education.
Moallemi, Ciamac C., Beomsoo Park, and Benjamin Van Roy (2012). "Strategic Execution in the Presence of an Uninformed Arbitrageur". In: *Journal of Financial Markets* 15.4, pp. 361–391.
Modigliani, Franco and Merton H Miller (1958). "The cost of capital, corporation finance and the theory of investment". In: *The American economic review* 48.3, pp. 261–297.
Moore, Tyler and Nicolas Christin (2013). "Beware the middleman: Empirical analysis of Bitcoin-exchange risk". In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 25–33.
Moore, Tyler, Nicolas Christin, and Janos Szurdi (2018). "Revisiting the Risks of Bitcoin Currency Exchange Closure". In: *ACM Transactions on Internet Technology* 18.4, 50:1–50:18.
Möser, Malte and Rainer Böhme (2015). "Trends, tips, tolls: A longitudinal study of Bitcoin transaction fees". In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 19–33.
Möser, Malte, Rainer Böhme, and Dominic Breuker (2013). "An inquiry into money laundering tools in the Bitcoin ecosystem". In: *2013 APWG eCrime researchers summit*. Ieee, pp. 1–14.

Mossin, Jan (1966). “Equilibrium in a Capital Asset Market”. In: *Econometrica* 34.4, pp. 768–783.
Nadarajah, Saralees and Jeffrey Chu (2017). “On the inefficiency of Bitcoin”. In: *Economics Letters* 150, pp. 6–9.
Nakamoto, Satoshi (2008a). “Bitcoin: A peer-to-peer electronic cash system”. In:
— (2008b). *Re: Bitcoin P2P e-cash paper*. Mail archive. Online at `https://bit.ly/2HMIn6L`. Accessed on 29 October 2020.
Nan, Zheng and Taisei Kaizoji (2019). “Bitcoin-based triangular arbitrage with the Euro/US dollar as a foreign futures hedge: modeling with a bivariate GARCH model”. In: *Quantitative Finance and Economics* 3.2, pp. 347–365.
Narayanan, Arvind et al. (2016). *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press.
Newman, Mark EJ (2005). “Power laws, Pareto distributions and Zipf’s law”. In: *Contemporary Physics* 46.5, pp. 323–351.
Petrov, Andreas and David Schufla (2013). “Arbitrage am Bitcoin-Markt”. In: *1st Münster Bitcoin Conference*. Ed. by Rainer Böhme and Dominic Breuker. University of Münster.
Pichl, Lukáš and Taisei Kaizoji (2017). “Volatility Analysis of Bitcoin Price Time Series”. In: *Quantitative Finance and Economics* 1.4, pp. 474–485.
Pieters, Gina and Sofia Vivanco (2015). *Bitcoin arbitrage and unofficial exchange rates*. Tech. rep. Trinity University.
— (2017). “Financial regulations and price inconsistencies across Bitcoin markets”. In: *Information Economics and Policy* 39, pp. 1–14.
Preneel, Bart (2010). “The first 30 years of cryptographic hash functions and the NIST SHA-3 competition”. In: *Cryptographers’ track at the RSA conference*. Springer, pp. 1–14.
Reynolds, Julia et al. (2018). “Deviations from Triangular Arbitrage Parity in Foreign Exchange and Bitcoin Markets”. In:
Rivest, Ronald L, Adi Shamir, and Leonard Adleman (1978). “A method for obtaining digital signatures and public-key cryptosystems”. In: *Communications of the ACM* 21.2, pp. 120–126.
Roll, Richard (1984). “A simple implicit measure of the effective bid-ask spread in an efficient market”. In: *The Journal of finance* 39.4, pp. 1127–1139.
Roll, Richard and Stephen Ross (1980). “An empirical investigation of the arbitrage pricing theory”. In: *The Journal of Finance* 35.5, pp. 1073–1103.

Ron, Dorit and Adi Shamir (2013). “Quantitative analysis of the full bitcoin transaction graph”. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 6–24.
Ross, Stephen (1976). “Return, risk and arbitrage”. In: *Risk and return in finance* 1, pp. 189–218.
— (1978). “A simple approach to the valuation of risky streams”. In: *Journal of business*, pp. 453–475.
Ross, Stephen A (1976). “The arbitrage theory of capital asset pricing”. In: *Journal of Economic Theory* 13.3, pp. 341–360.
Ruiz, Jesus (2020). *Public-Permissioned blockchains as Common-Pool Resources*. Tech. rep. Alastria Blockchain Ecosystem.
Scaillet, Olivier, Adrien Treccani, and Christopher Trevisan (2017). *High-frequency jump analysis of the bitcoin market*. Tech. rep. Swiss Finance Institute, Research Paper Series, No. 17-19.
Shafir, Eldar, Peter Diamond, and Amos Tversky (1997). “Money illusion”. In: *The Quarterly Journal of Economics* 112.2, pp. 341–374.
Sharpe, William F. (1964). “Capital Asset Prices — A Theory of Market Equilibrium Under Conditions of Risk”. In: *Journal of Finance* 19.3, pp. 425–442.
Shleifer, Andrei and Robert W. Vishny (1997). “The Limits of Arbitrage”. In: *Journal of Finance* 52.1, pp. 35–54.
Shynkevich, Andrei (2020). “Bitcoin Arbitrage”. In: *Finance Research Letters*, pp. 293–319.
Smith, Jacob (2016). “An analysis of bitcoin exchange rates”. In:
Stoll, Hans R (2000). “Friction”. In: *The Journal of Finance* 55.4, pp. 1479–1514.
Szabo, Nick (2005). *Bit gold*. Website-Blog. See `https://nakamotoinstitute.org/bit-gold/`.
Thode, Henry C (2002). *Testing for normality*. Vol. 164. CRC press.
Tiwari, Aviral Kumar et al. (2018). “Informational efficiency of Bitcoin—An extension”. In: *Economics Letters* 163, pp. 106–109.
Underwood, Barbara (Sept. 2018). *Virtual Markets Integrity Initiative*. Office of the New York State Attorney General.
Urquhart, Andrew (2016). “The inefficiency of Bitcoin”. In: *Economics Letters* 148, pp. 80–82.
Varian, Hal R (1987). “The arbitrage principle in financial economics”. In: *Journal of Economic Perspectives* 1.2, pp. 55–72.
Vasek, Marie and Tyler Moore (2015). “There’s no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency

scams". In: *International conference on financial cryptography and data security*. Springer, pp. 44–61.
Wang, Ye et al. (2021). "Cyclic Arbitrage in Decentralized Exchange Markets". In: *Available at SSRN 3834535*.
Wheatley, Spencer et al. (2018). *Are Bitcoin Bubbles Predictable? Combining a Generalized Metcalfe's Law and the LPPLS Model*. Tech. rep. Swiss Finance Institute, Research Paper Series, No. 18-22.
Wood, Robert A, Thomas H McInish, and J Keith Ord (1985). "An investigation of transactions data for NYSE stocks". In: *The Journal of Finance* 40.3, pp. 723–739.
Yermack, David (2015). "Is Bitcoin a real currency? An economic appraisal". In: *Handbook of Digital Currency*. Elsevier, pp. 31–43.
Yin, Haohua Sun and Ravi Vatrapu (2017). "A first estimation of the proportion of cybercriminal entities in the bitcoin ecosystem using supervised machine learning". In: *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, pp. 3690–3699.
Yu, Gloria and Jinyuan Zhang (2018). "A Revisit to Capital Control Policies When Bitcoin Is in Town". In: *31st Australasian Finance and Banking Conference*.