



Arbitrageurs in the Bitcoin ecosystem: Evidence from user-level trading patterns in the Mt. Gox exchange platform

Pietro Saggese^{a,b,c}, Alessandro Belmonte^{d,e,*}, Nicola Dimitri^f, Angelo Facchini^a, Rainer Böhme^g

^a IMT School for Advanced Studies Lucca, Italy

^b AIT - Austrian Institute of Technology, Austria

^c Complexity Science Hub Vienna, Austria

^d University of Urbino, Italy

^e CAGE, University of Warwick, UK

^f Università di Siena, Italy

^g Department of Computer Science, Universität Innsbruck, Austria

ARTICLE INFO

JEL classification:

C58
D53
G11
G40

Keywords:

Arbitrage
Bitcoin
Cryptocurrency exchanges
Behavioral finance

ABSTRACT

We mine the leaked history of trades on Mt. Gox, the dominant Bitcoin exchange from 2011 to early 2014, in order to detect the triangular arbitrage conducted on the platform. To this end, we exploit user identifiers per trade to identify and describe the individual trading patterns of 440 arbitrageurs. Moreover, we introduce proxies for expertise and document that the expert users' distribution of profits first-order stochastically dominates that of non-expert users. Most importantly, by including user fixed effects, we show that expert users make profits on arbitrage by reacting quickly to plausible exogenous variations on the official exchange rates. A small number of expert arbitrageurs are able to conduct the vast majority of the arbitrage actions and systematically yield higher profits: our results provide empirical evidence that arbitrageurs are few and sophisticated users, characterized by the ability to incorporate information and to quickly react to exogenous shocks within short time scale intervals.

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 vitally 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 (Ross, 1976), the theory of option pricing (Merton, 1973; Black and Scholes, 1973), the efficient market hypothesis (Fama, 1970), and many other theories.

* Corresponding author.

E-mail addresses: pietro.saggese@imtlucca.it (P. Saggese), alessandro.belmonte@uniurb.it (A. Belmonte), nicola.dimitri@unisi.it (N. Dimitri), angelo.facchini@imtlucca.it (A. Facchini), rainer.boehme@uibk.ac.at (R. Böhme).

<https://doi.org/10.1016/j.jebo.2023.07.025>

Received 6 December 2022; Received in revised form 18 May 2023; Accepted 17 July 2023

Available online 28 July 2023

0167-2681/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

In practice, arbitrage is never risk-free. Since purchases and sales are not executed in an atomic¹ 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 have the knowledge and information to engage in arbitrage.” (p. 36) While this is plausible and has likely been cross-checked by expert market participants, the evidence remains anecdotal. In this work, we use rich micro data at the trader level to provide a partial answer to this question, limiting our scope to investigating whether the arbitrageurs are many small traders or a few sophisticated investors.

We mine a leaked dataset of individual and identified trades from Mt. Gox, a now-defunct exchange between convertible currency and cryptocurrency that enjoyed a dominant market position in the early years of Bitcoin, before its collapse in 2014. Crucially for our purposes, Mt. Gox allowed users to trade within the same exchange bitcoins against different fiat currencies, thus providing opportunities to execute triangular arbitrage activities. Using bitcoin as a vehicle currency, investors could compare the implied relative price of traditional currencies (information that we observe in the private ledger of a single exchange) to the official exchange rate and look for the presence of mispricings. Moreover, the exchange limited users to only one personal account at a time to which a unique label was assigned. We exploit these two unique features of Mt. Gox to identify the arbitrage actions through an algorithm based on the analysis of identified sequences of trades conducted by the same user. We detect 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.²

After having identified which trading actions constitute arbitrage, and therefore which players act as arbitrageurs in the exchange, we quantify the magnitude of the triangular arbitrage activity within the Mt. Gox platform. Consistent with the anecdotal evidence, we find that only a restricted group of users perform at least one arbitrage action, while an even smaller group of sophisticated users are responsible for the vast majority of trades; the users in this subset tend to be active in multiple currency markets rather than in a single one, they conduct complex strategies (i.e., metaorders), and prefer limit to market orders.³ We then introduce proxy measures for the trading ability of the investors to determine whether expert users conduct more profitable actions. We first note that the expert users' distribution of profits first-order stochastically dominates that of non-expert users. Next, using variation across trades executed using the same pair of currencies and within the same hour, and controlling for the volume of the trades, expressed in USD dollars, we estimate an average profit achieved by the expert arbitrageurs that is 1.29% (of the hourly official rate between the fiat currencies) higher than that obtained by the unsophisticated arbitrageurs — a difference which is slightly above a standard deviation in profitability. The arbitrage activity of sophisticated investors is on average profitable. Instead, the arbitrage activity attributable to the non-expert users is, to a large extent, non-profitable. These findings hold for different definitions of expertise and are unlikely to be explained by a “learning-by-doing” process or by the specific way in which we define a triangular arbitrage action.

Next, we investigate why expert users are more likely to make profits on arbitrage relative to non-experts. Using a wide range of definitions for trader expertise, we document that expert arbitrageurs are more able to exploit temporary arbitrage opportunities in their favor. Specifically, we use the (unsigned) rate of change of the official exchange rate between two fiat currencies to capture such temporary arbitrage opportunities and demonstrate that expert users make profits on arbitrage by reacting quickly to plausible exogenous variations on the official exchange rates. Remarkably, this finding also holds when we include user fixed effects, which allow us to absorb a relatively large set of unobservables at the user level that are likely to correlate with our measure of trade ability, including education, financial literacy, and wealth. Varying the way in which we identify a triangular arbitrage action yields similar estimate results. Consistent with our story, we find that arbitrage profits increase (and so too does the premium the expert arbitrageurs attain relative to the non-sophisticated traders) when we restrict the sample of triangular arbitrage to the actions completed during available price deviations. Finally, we also provide several falsification tests that help to rule out “learning” as a likely channel for our findings. Overall, these findings indicate that arbitrageurs are few and sophisticated users, characterized by the ability to incorporate information and quickly react to exogenous shocks within small time scale intervals. While these results are drawn from the analysis of a peculiar cryptocurrency market, we exhaustively discuss below how to interpret them if applied to other temporal and market contexts.

Our study contributes to the understanding of arbitrageurs' behavior by examining their investment strategies as well as their responsiveness to temporary opportunities that arise from their comparisons of different market indices. In this respect, our paper

¹ Incidentally, this has changed with the introduction of decentralized exchanges on programmable cryptocurrency platforms (Makridis et al., 2023).

² A similar approach to ours has been proposed by Luckner et al. (2023) to identify cross-border Bitcoin capital flows and by Aloosh and Li (2019) to detect wash trading. The latter exploit the same dataset, but their algorithm design detects a wash trade when its buy and sell legs have the same user ID; thus, the two approaches identify non-overlapping sets of trades.

³ This finding is consistent with Westphal and Sornette (2020) model, which describes heterogeneous traders engaging in different arbitrage strategies.

complements prior works (e.g., Roll and Ross, 1980; Malkiel, 2003; Lamont and Thaler, 2003) that assess the validity of consolidated theories on arbitrage based on the analysis of aggregate information on trading in several markets. We build on these works by providing econometric evidence that is inferred from micro-data, that is, from the analysis of the behavior of individual traders over a two-year period. In doing this, our study also relates to recent works that use financial data at the trader level. Hasso et al. (2019), for example, exploit user-level data from an online brokerage service to investigate the individual characteristics of cryptocurrency investors, such as their demographics and other trading activities, but they do not investigate individual sequences of trades. Meanwhile, other studies that use user-level data lack identifiers (e.g., Lee and Ready, 1991; Lee and Radhakrishna, 2000), or focus largely on risk profiling (e.g., Kourtidis et al. (2011); De Bortoli et al. (2019); Borsboom and Zeisberger (2020)). Exchanges typically guarantee anonymity to their customers, and distributed ledger technologies are publicly auditable, but possibly privacy-preserving, making it difficult to track the same user over time. A prominent example of this is represented by Wang et al. (2022), who identify the cyclic arbitrages executed in decentralized exchanges (DEXs) within the Ethereum ecosystem. However, the Ethereum protocol design limits the potential to trace users exactly.

Furthermore, our results speak to the literature that investigates market anomalies despite the presence of rational investors in financial markets (Harris and Gurel, 1986; De Long et al., 1990; Froot and Dabora, 1999; Lamont and Thaler, 2003).⁴ Our paper partially fills an existing gap between the theoretical description of the arbitrage activity and the practical evidence from real markets by providing empirical evidence that arbitrageurs are few and sophisticated in trading (Shleifer and Vishny, 1997). What is more, we identify a group of non-sophisticated arbitrageurs who perform arbitrage in a non-profitable way — a fact that we demonstrate is reliant on their poor ability to quickly exploit opportunities in the market as well as to take into account transaction costs in their arbitrage decisions (in a mechanism that is similar to the monetary illusion phenomenon; see Shafir et al. (1997)).⁵

2. Background

2.1. Cryptocurrency exchanges and arbitrage strategies

Bitcoin, the most prominent cryptocurrency in terms of market capitalization,⁶ 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 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 where price formation occurs (Halaburda et al., 2022). 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.

Two arbitrage strategies are particularly relevant for cryptocurrency markets: arbitrage across exchanges and within one exchange. In the former, arbitrageurs maintain a stock of both bitcoins and fiat money in accounts at multiple exchanges to enable them to react quickly to price differences. The funds can be balanced at a lower frequency and are not necessarily correlated with observable price differences. Therefore, while arbitrage opportunities are measurable from published data, it is more difficult to identify arbitrageurs from the public ledger. The latter 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. We restrict our analysis to the second strategy, as this form of arbitrage has the advantage that information required to identify the two legs composing one arbitrage action is contained in the private ledger of a single exchange (additional external information is needed to compare the implied exchange rates across fiat currencies within the market to the prices outside of the market).

2.2. Literature on triangular arbitrage in cryptocurrency markets

Triangular arbitrage in the Bitcoin ecosystem has been widely investigated in the literature, documenting systematic unexploited arbitrage opportunities in cryptocurrency markets, especially before 2018. Focusing on the early years of Bitcoin trading (which are relevant for our work), Dong and Dong (2015) test for the presence of triangular arbitrage opportunities between the main cryptocurrency exchanges and the spot currency markets, finding evidence of persistent price deviations. Similarly, Smith (2016) examines Mt. Gox aggregate data, finding that shocks in that market did not affect rates in conventional venues and that the efficiency observed in the market could be explained by the presence of arbitrageurs. Other studies indicate the presence of triangular arbitrage opportunities during the period from 2013 to 2017 (Pichl and Kaizoji, 2017; Reynolds et al., 2021; Pieters and Vivanco, 2015; Makarov and Schoar, 2020; Yu and Zhang, 2018; Hirano et al., 2018; Nan and Kaizoji, 2019). Remarkably, Reynolds et al. (2021) show that persistent mispricings arise only when Bitcoin is used as a vehicle currency, finding no evidence of deviations from parity when considering the implied rate between traditional fiat currencies. Makarov and Schoar (2020) and Yu and Zhang (2018) both

⁴ For a survey on the limits of arbitrage, see Barberis and Thaler (2003) and Gromb and Vayanos (2010).

⁵ Finally, in describing the Mt. Gox market structure and its internal trading mechanisms, we also make a connection to the literature on the role of cryptocurrency exchanges in the Bitcoin ecosystem (e.g., Moore et al., 2018; Griffin and Shams, 2020). In a similar vein to Dyhrberg et al. (2018), who investigate transaction costs and the liquidity of bitcoin, we account for explicit transaction costs paid within a cryptocurrency exchange (in contrast to previous research that focuses on the transactions costs in the Bitcoin network; see e.g., Möser and Böhme (2015), Dimitri (2019), and Easley et al. (2019)).

⁶ Bitcoin is valued at around 500 billion \$ at the time of writing (according to Coinmarketcap.com <https://bit.ly/3iXhZnj>).

indicate that capital controls are behind market frictions that give rise to arbitrage opportunities, while Hirano et al. (2018) show that such opportunities are more frequent in minor currency markets. However, recent studies suggest that in more mature markets, and especially after 2018, price deviations scarcely exist because of the presence of informed institutional investors and websites such as bitgap, tokenspread, and cryptohopper that collect information on mispricings before providing it to retail investors (Borri and Shakhnov, 2022; Crépellière et al., 2023).

In summary, based on heterogeneous methods and the study of periods with data of different frequency, the literature relatively consistently reports unexploited arbitrage opportunities in cryptocurrency markets, especially before 2018. This does not imply that arbitrage did not occur but rather that the costs and risks of arbitrageurs are underestimated. Indeed, anecdotal evidence from forums, the existence of web-based arbitrage tools, and code repositories for trading bots indicate that arbitrage did occur even at the early stages of the Bitcoin ecosystem (Petrov and Schufli, 2013).

2.3. The Mt. Gox exchange and the leaked dataset

All the studies reviewed above have in common that they analyze aggregated price (and sometimes volume) time series. Our approach represents a departure from these studies in that we use individual-level data from the internal ledger of a major exchange, Mt. Gox.

Mt. Gox played a prominent role during the early years of Bitcoin. It was established in 2010 as essentially the first cryptocurrency exchange and dominated the market with around 80–90% of the total trading volume until late 2012. Moreover, 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, and BTC China at the end of 2011.

Since the beginning of Spring 2013, a series of events gradually undermined the credibility of Mt. Gox,⁷ with customers starting to experience delays when withdrawing fiat money.⁸ Consequently, the volume traded in Mt. Gox decreased significantly in the following months, and the bitcoins started to be traded at a large premium in Mt. Gox⁹: 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. The exchange stopped withdrawals at once on 7 February 2014, and filed for bankruptcy two weeks later. The former CEO was then arrested after being charged with 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 et al. (2018).

Our main dataset was leaked to the public in 2014 as a series of CSV files. They contain around 7.5 million trades executed from 1 April 2011 to 30 November 2013. Each of these trades is composed of two buy and sell legs, reported in separate rows. Some variables are 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 Online Appendix Section C. The vast majority (87.9%) of trades are in USD, followed by EUR (7.7%).

Fig. 1 visualizes selected indicators on how Mt. Gox's user base has evolved over time, reaching a total of more than 125,000 by the end of 2013. The plots outline intuitively that the peaks of interest towards the cryptoasset (Panel a) and of activity within the market (Panel b) correspond with periods of exponential growth of the bitcoin price.

Whilst the dataset covers a longer period of time, we restrict our sample and exclude the trades executed after March 2013 due to the increasing difficulties in managing withdrawal operations and because of the other factors documented above. The total number of active users from April 2011 to March 2013 was approximately 72,000, while the number of trades were around 5.5 million. We further restricted the sample because the option to trade in currencies other than USD was introduced in September 2011.¹⁰ Thus, the identified arbitrage actions fall in the time period ranging from September 2011 to March 2013. As Fig. 1 shows, the time window considered also coincides with an epoch of constant activity within the exchange platform as well as with the linear growth rate of new registered users.

The leaked dataset has been widely analyzed already by a number of prior works but in relation to research topics distinct from arbitrage, e.g., the presence of metaorder executions (Donier and Bonart, 2015), unusual price jumps in the BTC/USD exchange rate (Scaillet et al., 2020), the effects of distributed denial-of-service (DDoS) attacks on trading activity (Feder et al., 2018), the impact of suspicious activity in the Mt. Gox exchange that likely engaged price manipulation (Gandal et al., 2018, and Chen et al., 2019, the latter trying to answer the same question through the lenses of network science), herding behavior (Haryanto et al., 2019), and wash trading (Aloosh and Li, 2019)).

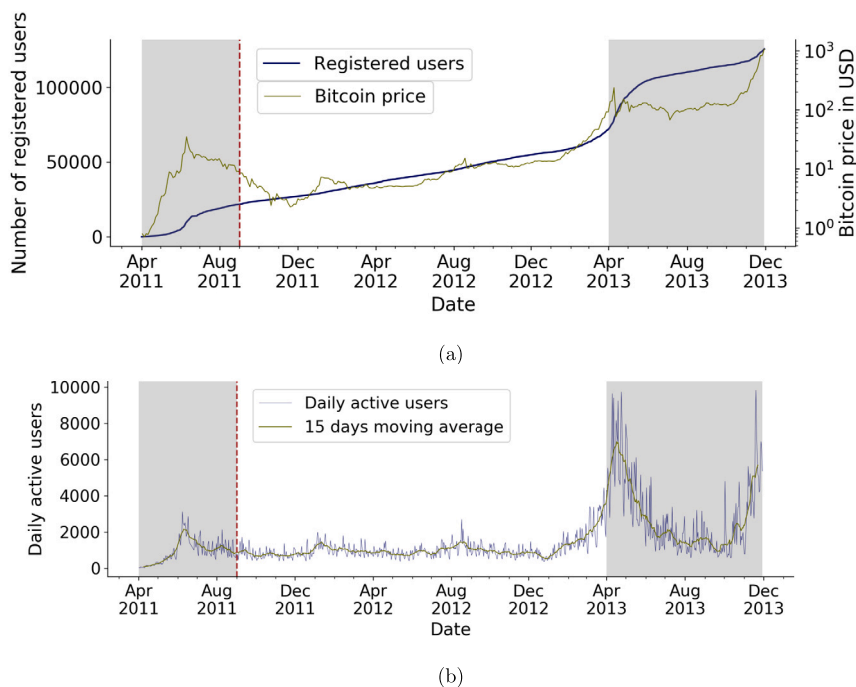
In summary, the dataset we explore in this study has been analyzed by various previous researchers and is therefore already largely accepted in the literature. This, along with our own comparisons to external aggregate information reported in Online

⁷ 11 March 2013: Mt. Gox suspends bitcoin deposits after hard fork. <https://bit.ly/2GWpKj>; 11 April 2013: Mt. Gox went down after unexpected increase in the trading activity. <https://bit.ly/3IEuSUW>; 2 May 2013: Coinlab files a lawsuit against Mt. Gox <https://bit.ly/3duI755>; 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>.

⁸ 18 April 2013: users point out withdrawal delays. <https://bit.ly/3IHAYO7>; 4 July 2013: Mt. Gox resumes the U.S. withdrawals halted on 20 June. <https://bit.ly/3doQxxH>, <https://bit.ly/3IEvPFY>.

⁹ See <https://bit.ly/2FshVy6>.

¹⁰ From September 2011 onwards, users were also permitted to trade bitcoins in exchange for EUR, CAD, GBP, CHF, RUB, AUD, SEK, DKK, HKD, PLN, CNY, SGD, TBH, NZD, and JPY (<https://bit.ly/314a5Cg>).



Notes: Panel (a) shows the growth of registered users in relation to the bitcoin price (the latter is reported 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 sample.

Fig. 1. Descriptive statistics of Mt. Gox users.

Appendix Section C, support its validity and authenticity.¹¹ Moreover, according to the *Guardian*,¹² several members of the Bitcoin community claimed to have found their own transactions in the dataset. Finally, certain facts established in the court case against the former CEO of Mt. Gox seem to plausibly explain patterns in the dataset.¹³

We rely on the work of Gandal et al. (2018), Feder et al. (2018), and Scaillet et al. (2020) to pre-process and clean the original leaked dataset. This stage chiefly consists of finding duplicate rows and identifying misreported data. The procedure is described in more detail in Online Appendix Section C. It is worth noting that our aggregation technique differs from the above reference in that we aggregate the trades belonging to the same user occurring within the same second. Put differently, we assume that such actions belong to the same executed order, in compliance with the operating principle of the Mt. Gox filling mechanism.¹⁴ Order speed analyses on other cryptocurrency exchanges reveal that a one-second time scale is suitable for measuring order execution delays¹⁵; traditional financial markets show much shorter latencies (see, e.g., Budish et al. (2015); Hasbrouck and Saar (2013); 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 would become public. We therefore take the utmost care in ensuring that none of our analyses singles out users that have not been singled out in other work. Furthermore, user identifiers in our figures cannot be directly related to identifiers in the data source. Therefore, we believe that the harm caused by our study is minimal, whereas the benefits that we provide by shedding light on an important area of finance are very clear.

3. Identification and description of arbitrage activity

3.1. Detection of the 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 our context, bitcoin) diverges from the official exchange rate. Thus, an investor

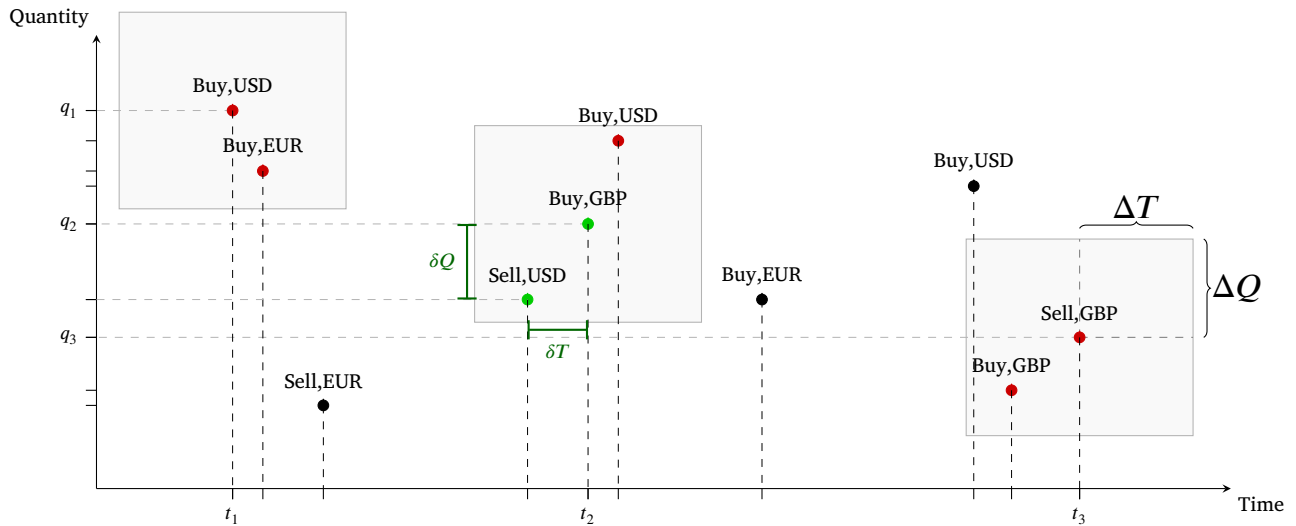
¹¹ Comparisons refer to the data published by Bitcoincharts.com. In addition, as in Scaillet et al. (2020), we match our data with an aggregated dataset published by Mt. Gox.

¹² <https://bit.ly/2lu77Rk>.

¹³ This statement is based on personal communication. The authors have not read the Japanese files.

¹⁴ <https://bit.ly/33YCxaG>.

¹⁵ 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 10 ms to 100/200 ms, with skewed tails, while there are a non-negligible number of trades whose latency approaches the second level.



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: we exclude the legs in the neighborhood of $[t_1, q_1]$, as 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.

Fig. 2. Algorithm to identify the arbitrage actions. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

seeking to exploit such opportunities requires access to at least two currency markets quoted against the same third currency: Mt. Gox users could trade within the same platform in multiple fiat-to-bitcoin markets and were entitled to have only one personal account at a time.¹⁶ Bearing in mind that all the legs are labeled by individual identifiers, this setting is ideal to study triangular arbitrage at the micro (individual) level. We start by selecting only those investors who traded bitcoins against more than one fiat currency, which amounts to only 3,825, out of 71,808 users in our sample; nonetheless, around 1,600,000 legs in the leaked dataset are attributable to them. A subgroup of 307 investors exchanged bitcoins for more than two fiat currencies, being involved in around 800,000 legs.

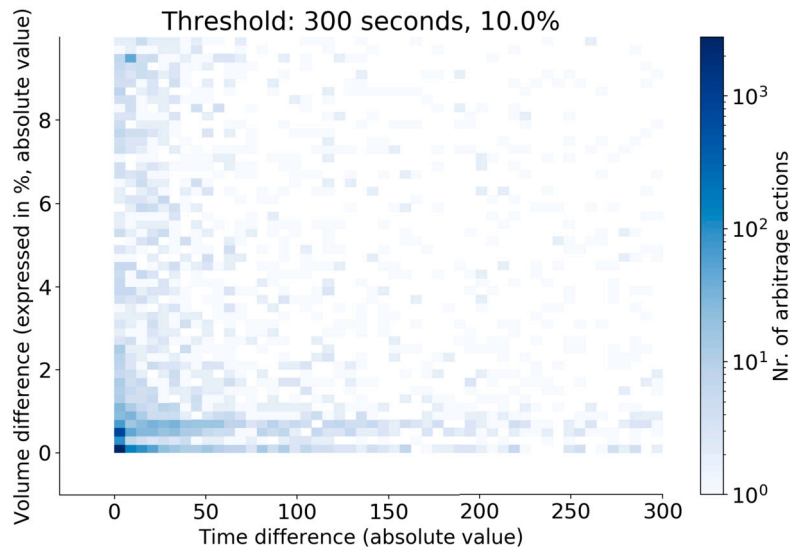
Next, we implement an algorithm to identify arbitrage actions. The underlying idea of the algorithm is illustrated in Fig. 2. Using the information in the leaked dataset, it is possible to detect a potential arbitrage action in the form of a pair of (buy, sell) legs executed in two separate trades, by the same investor, and using different currencies, such that the first leg closes the second leg. Textbook definitions posit that arbitrage is performed through simultaneous actions involving equivalent securities. Our data, however, highlight a mass of potentially triangular actions that are temporally retarded and/or executed with a marginal volume mismatch. We illustrate the distribution of these potentially triangular actions in Fig. 3 in the space $(\delta T, \delta Q)$, respectively the distance in time and volume between the two legs composing such action. While the distribution presents, as expected, a marked peak in the density in the proximity of the origin (i.e. $\delta T = 0$ s, $\delta Q = 0\%$), we note a small-scale probability mass around the origin.

To ensure that we do not discard relevant information, we account for the possibility of a marginally imperfect match and therefore set reasonably small boundaries for the maximum time delay, ΔT , and volume difference, ΔQ , between the pair of legs, that represent the two parameters of the algorithm¹⁷; likewise, to ensure that we are not including false positive actions, we provide alternative identification strategies varying both ΔT and ΔQ (and especially including $\Delta Q = 0\%$) and, furthermore, by restricting the sample to actions that are traded when a likely arbitrage opportunity occurs (i.e., when the official exchange rate exceeds a certain threshold). Reassuringly, all these robustness checks, which will be presented below, are consistent with each other and provide evidence that our conclusions do not depend on the choice of the specific values of the parameters ΔT and ΔQ .

We therefore explore the Mt. Gox log searching for pairs of buy and sell legs that move a nearly equivalent amount of bitcoin, 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. In our main analysis, we identify triangular arbitrage actions that are executed within a maximum delay of up to 300 seconds and with a maximum volume mismatch of 10% (that is, using the boundaries $\Delta T = 300$ s and $\Delta Q = 10\%$). The resulting sample comprises 6,629 actions.

¹⁶ <https://bit.ly/2Fu1Rfp> and <https://bit.ly/3j0EkAl>.

¹⁷ The parameters ΔT and ΔQ and the action-specific variables δT and δQ are defined in Table 1.



Notes: each arbitrage action is characterized by a δT and a δQ , representing the distance in time and volume between the two legs composing such action. By construction, they are respectively smaller than $\Delta T = 300$ s and $\Delta Q = 10\%$. Darker shades indicate a higher number of actions (in logarithmic scale).

Fig. 3. Distribution of the arbitrage actions by δT and δQ , given ΔT and ΔQ .

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. We then measure the profitability of an arbitrage action as follows:

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

where $OffER$ is the hourly official rate¹⁸ and $ImpER$ is the implied one. To compare them, by construction we use the direct quotation with the currency of the buy leg acting as the (fixed) foreign currency. That is,

$$OffER = CUR_B to CUR_S,$$

where CUR_B is the fiat currency used to trade bitcoins on the buy side, and CUR_S on the sell side,¹⁹ and

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

Notably, the leaked log includes information on the explicit transaction costs incurred by the Mt. Gox users (i.e., the fees associated with each leg of all trades).²⁰ Although additional costs may (and are in fact likely to) exist,²¹ this feature of the dataset is especially important, as it allows us to account for the costs within the exchange that directly affect the profitability of the arbitrage activity. Thus, in the baseline investigation we adjust the actual profitability by incorporating the leg-specific fees in the prices paid to trade bitcoins, as described formally in Table 1, which provides a recap of the main variables introduced in this work. However, as a robustness check, we consider two additional ways to account for the explicit fees (i.e., in the first scenario we exclude them, in the other scenario we estimate the fees a user would expect to pay given the official Mt. Gox schedule), which are discussed in Online Appendix Section B.

¹⁸ Our choice of using the hour-level official exchange rate, instead of a more granular scale, relies on two considerations: First, we observe little variation in the main exchange rates at a lower level, which would have considerably undermined the statistical power of our analysis. Second, we allow an investor to complete a triangular arbitrage action up to 300 seconds. For consistency, we need to measure profits at a time interval that is sufficiently large to carry all the relevant information to the trader in taking his/her investment decision.

¹⁹ We use the hourly open prices of the official exchange rates published on <https://www.histdata.com/>. This dataset lacks information for a few minor currencies (CNY, THB, NOK, RUB). As a result, we 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 Fig. 5d conducted 796 arbitrage actions, but we can calculate the profitability only for 782 of them.

²⁰ 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>.

²¹ In addition to the Mt. Gox fees, investors potentially also incur implicit transaction costs and other explicit costs when closing the triangular arbitrage action on external markets, e.g., trading fees or commissions, bid-ask spreads due to low liquidity, or concurrent price changes (Stoll, 2000).

Table 1
Definitions of the main variables.

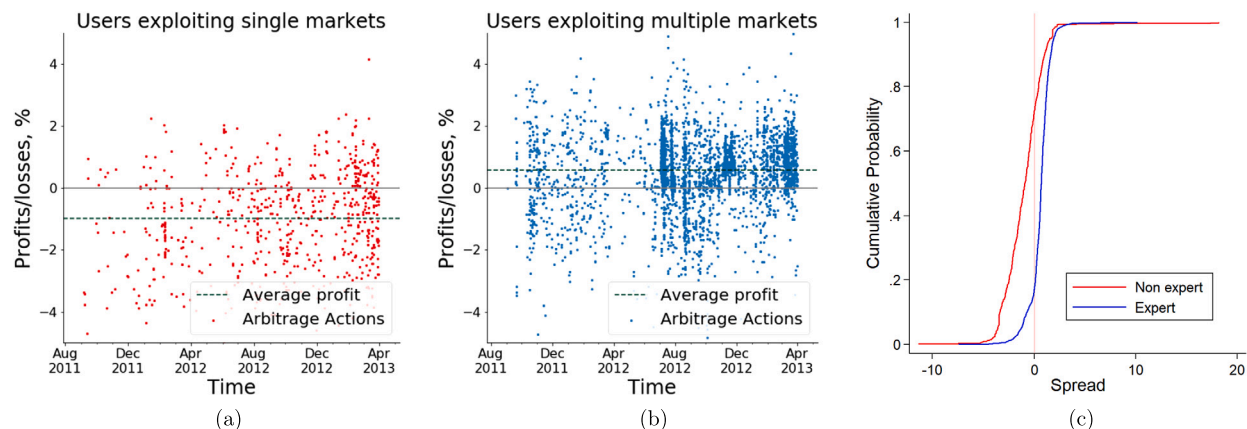
Variable	Description/formalization
Arbitrage action	Action composed of two legs executed by the same user in different trades using different currencies. The time delay and volume difference cannot exceed a threshold $[\Delta T, \Delta Q]$. Each arbitrage action is obtained by merging a buy and sell leg. $ArbitrageAction = (Leg_{Buy}, Leg_{Sell}) \tag{1}$
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 : $\delta T = T_B - T_S \leq \Delta T \tag{2}$
δQ	Volume difference (Bitcoins traded) between Leg_B and Leg_S , expressed as a percentage. By definition smaller or equal to ΔQ : $\delta Q = \frac{ Vol_B - Vol_S }{(Vol_B + Vol_S)/2} \cdot 100 \leq \Delta Q \tag{3}$
Official exchange rate	By convention, each arbitrage action is compared to the official exchange rate in the following way: $OfferER = CUR_B to CUR_S \tag{4}$ that is, if the buy leg of an arbitrage action is performed in EUR and the Sell one is in USD, then we 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.
$ \Delta R $	Hourly unsigned percentage variation of the official exchange: $ \Delta R = \frac{ OfferER_{t_1} - OfferER_{t_0} }{OfferER_{t_0}} \cdot 100 \tag{5}$
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. $ImpER = \begin{cases} \frac{Fiat_S}{BTC_S} \cdot \frac{BTC_B}{Fiat_B} & \text{without fees} \\ \frac{Fiat_S - Fee_{f,S}}{BTC_S + Fee_{b,S}} \cdot \frac{BTC_B + Fee_{b,B}}{Fiat_B - Fee_{f,B}} & \text{with fees} \end{cases} \tag{6}$
Profit (Spread)	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. Spread between the implied and the official rate divided by the official rate, as a percentage. By construction, profits arise when $ImpER > OfferER$. $Spread = \frac{ImpER - OfferER}{OfferER} \cdot 100 \tag{7}$
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. <i>Note:</i> we partly follow the methodology described in Donier and Bonart (2015), with some differences: the authors consider a larger time delta (one hour) between each action, and contrary to them we use an arbitrary parameter ($N=5$) to define the minimum length of a metaorder. While we do not provide the results here, we varied the two thresholds and noticed that the differences are negligible for our purposes.
Aggressive Equiv. \$	Arbitrage action composed by at least one aggressive leg (that is, a leg that initiated a market order). Value of a trade expressed in dollars. We use this variable to indicate the value of a trade since the bitcoin value is not stable in time.

The average arbitrage action is worthy of a profit which is 0.42% of the hourly official rate between the fiat currencies. The average amount of bitcoins traded are equivalent to 52 USD (see Panel A of Table 2 for summary statistics).

3.3. A preliminary inspection of the data

The structure of micro-data we collect allows us to uncover a number of patterns regarding the behavior and nature of the arbitrageurs. Notably, in disagreement with theory, we 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 users in our sample). Furthermore, the arbitrageurs' behavior seems to indicate 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

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 we will imply this concept.



Notes: the plot on the left reports the actions executed by arbitrageurs that exploited a single market. Viceversa, the central plot 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. We do not show them (here and in the following plots) to focus on the area of interest. Finally, the right panel shows the cumulative probability of the *Spread* variable for the expert (blue) and non-expert users (red). The former first-order stochastically dominates the latter.

Fig. 4. Profitability of the arbitrage actions. Users grouped by the number of currency markets exploited for arbitrage.

fiat-to-fiat markets by exchanging bitcoins for at least three fiat currencies. Significant differences are seen when comparing the two groups: Fig. 4 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 below the gray line are losses; the dashed line denotes 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 those in the former, which comprise traders who possess less expert knowledge, are on average negative. Even more importantly, Panel (c) shows that for any value of arbitrage profits the mass distribution for the second group (in blue) lies below that for the first group (in red). In other words, the former distribution first-order stochastically dominates the latter one. We therefore assume that this variable is a proxy of a user's level of expertise: investors who exploited a single market are less expert, while those that exploited multiple markets are more expert. In the following paragraphs we provide additional empirical evidence to support this assumption.

Panels B and C of Table 2 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 on average profitable. The expected fees 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. Differences in time distribution appear too. 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 fiat currency. To partially explain this unexpected result, we hypothesize that the conventional principal-agency relationship discussed by Shleifer and Vishny (1997) might not take place in this context. Indeed, we recall that the Bitcoin ecosystem was in its early stages at the time of our study, and large institutional investors did not engage in bitcoin trading (we will discuss the specific aspects of our context in Section 7).

Second, from Table 2, it also emerges 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, with the most active user performing just 27 actions. Table 3 provides further information on the number of actions executed by the two groups.

Third, arbitrageurs that operate on multiple markets are also more acquainted with sophisticated algorithms, such as metaorders. We 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 4, only 13 arbitrageurs executed metaorders, which are typically composed of fewer than 10 actions — each delayed by around 20 seconds — and moved an average amount of bitcoin equivalent to a few hundred dollars. Only five users performed more than five metaorders, all of whom exploited multiple currency markets and executed more than 100 arbitrage actions.

²³ Note that it is not our primary goal to precisely define metaorders; rather, we use this measure as an indicator to verify which users exploit these strategies more systematically. As a robustness check, we repeated the procedure by changing the minimum number of actions and the time delay, and found that the differences are negligible.

Table 2
Descriptive statistics of the arbitrage actions.

Panel A: all arbitrage actions (N = 6629)							
	Mean	St.D.	Min	25%	50%	75%	Max
Profits, fees, %	0.42	1.26	-11.35	0.064	0.618	1.094	18.16
P., exp. fees, %	0.28	1.22	-7.46	-0.186	0.358	0.979	18.24
P., no fees, %	1.04	1.21	-6.40	0.472	1.108	1.693	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, %	-0.98	1.94	-11.35	-2.167	-0.857	0.107	18.16
P., exp. fees, %	-0.93	1.90	-7.46	-2.151	-0.839	0.158	18.24
P., no fees, %	0.11	1.89	-6.40	-1.099	0.087	1.233	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	0.205	0.687	1.126	10.13
P., exp. fees, %	0.42	1.02	-7.34	-0.007	0.444	1.017	10.15
P., no fees, %	1.16	1.04	-6.28	0.573	1.174	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 $\Delta T = 300$ s and $\Delta Q = 10\%$. 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 3
Statistics on the number of actions executed by the arbitrageurs.

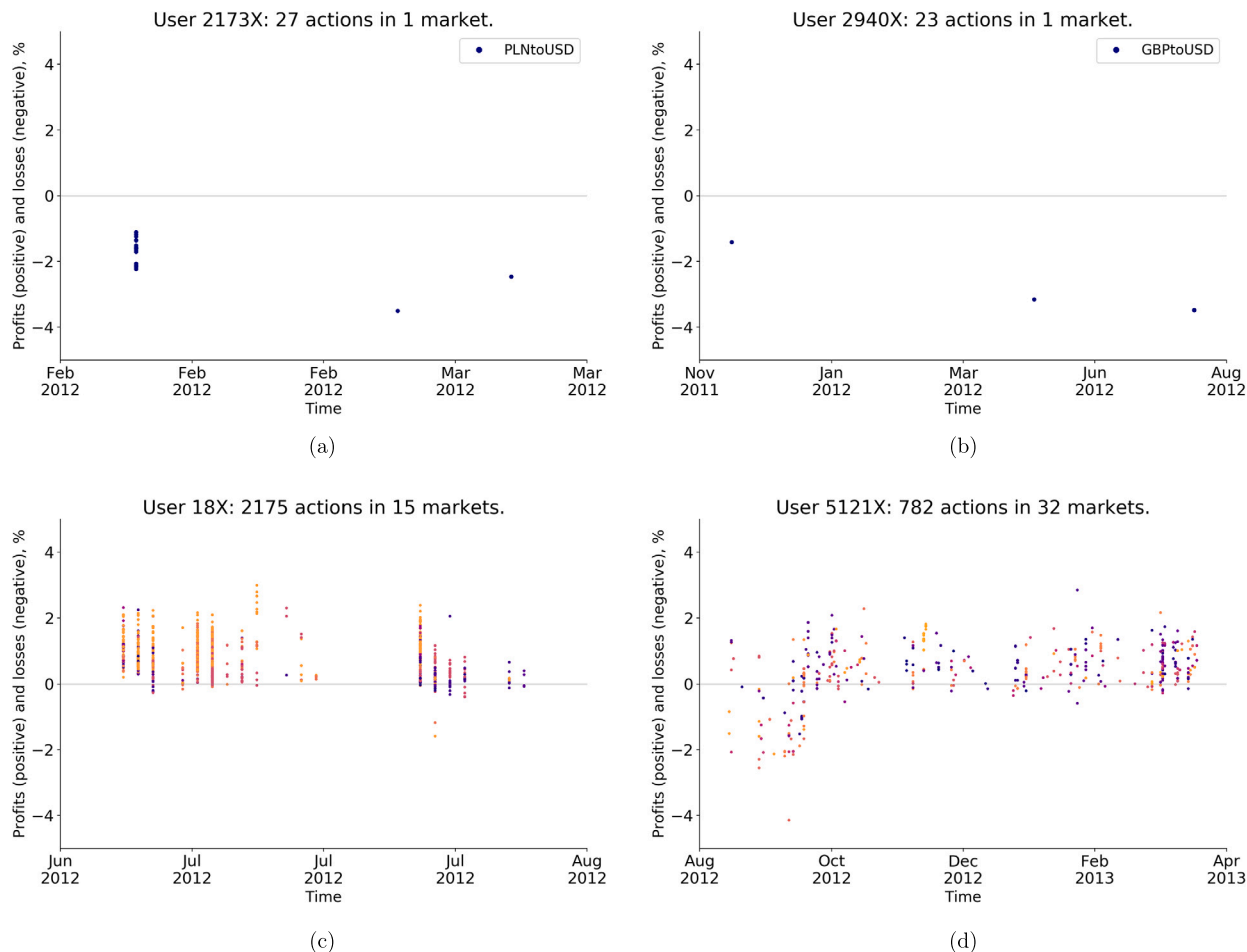
	Mean	Std	Min	25%	50%	75%	90%	95%	Max
Group <i>Single</i> (N = 395)	1	2	1	1	1	2	2	5	27
Group <i>Multiple</i> (N = 45)	131	366	2	4	11	28	392	690	2175

Notes: we 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 4
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), we 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 bitcoins and in dollars. User identifiers are anonymized.



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. We do not report the legend for the two plots below as the number of markets is too high (15 and 32). We 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.


Fig. 5. Profitability and trading patterns across arbitrageurs.

Table 5
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 (%)	-1.074	1.425	-5.354	-2.106	-0.911	0.0	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 (1st row: 6.57 on average, and maximum 28); on average they are not profitable (2nd row), and they are executed primarily by users active only on single markets (3rd row).

Fourth, arbitrageurs that operate on multiple markets are less likely to behave aggressively. In order book-based markets, traders that place bid and ask orders (i.e., limit orders) contribute to creating liquidity. Aggressive trades, though, demand liquidity from the order book by closing the limit orders. These are called market orders. Aggressive trades have a higher impact on costs: taking market orders usually incurs larger costs as liquidity is taken and not provided (Crépeillère et al., 2023). We follow Scaillet et al. (2020) 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 5 shows that aggressive orders have been used only by users who executed fewer than 30 actions and that, on average, they are not profitable.



Notes: the graph refers to arbitrageurs active on multiple markets. For about 70% of them, only 0 to 14 days passed between the first arbitrage action and the first one in another currency market (1st bin).

Fig. 6. Days passed between the first arbitrage action of a user and the first one in a new market.

Notably, clustering the arbitrage actions executed by the same users unearths interesting insights and provides further evidence that such differences map into heterogeneous patterns of profitability of arbitrage. In Fig. 5, for instance, we 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. From this, it can be seen that the differences in profits are considerable. While users in (a) and (b) systematically incur losses when trading (as dots lie below the gray line), the others typically make profits by executing far more complex trading patterns. It is also worth highlighting important differences between their strategies (e.g., when comparing users 18X and 5121X). Trades performed by the first group are concentrated in just a few weeks (around July 2012); these actions appear to be consequential and related and are likely to be part of one or a sequence of metaorders. The trading pattern of the second group, meanwhile, is steady and spans across a longer period of time. In spite of these differences, though, both strategies are profitable, non-trivial, and likely executed via algorithmic trading.

4. Trade ability and profitability of arbitrage

In this paper we hypothesize that arbitrage profitability is a function of the user's trade ability. As laid out in the previous section, our preferred indicator for trade ability is arbitrage on multiple markets (which we complement with three other variables — number of actions, execution of metaorders, and execution 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 conduct arbitrage in this form (and non-systematically, i.e., in few and dispersed trials, which is on average non-profitable). Few users, however, explore more than one market in search of 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 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). The descriptive evidence provided above suggests also that users that engage in arbitrage through multiple markets obtain higher profits. This conclusion is potentially threatened by two facts, though. First, trade ability may not be fixed but, rather, may increase with trading. Second, the correlation between trade ability and profits may be affected by an omitted variables bias. In this section, we take these two aspects into account in our analysis.

4.1. Learning-by-doing and trade ability

The validity of our analysis relies on the assumption that arbitrage through trading on multiple markets is a sign of trade ability that a user holds before they begin operating on Mt. Gox. Thus, our analysis fails to capture the link between expertise and profits if, for example, a user conducts arbitrage on a single market for an extended period of time and only after a period of training the user starts to conduct arbitrage using other currencies. In Fig. 6, we show that such a scenario is unlikely to hold in our sample. 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 can be seen, the distribution is concentrated in the first bin, which gathers arbitrageurs that operate on a new market within 14 days of its first arbitrage action. This bin collects approximately 70% of the arbitrageurs active in multiple markets, indicating that for the vast majority of arbitrageurs the time that passes between their first arbitrage action and the first one in a new currency market is short. While this is evidence that users' expertise does not change considerably over time, we note that the learning process could have occurred in an earlier period (i.e., before September 2011). For this reason, we conduct additional robustness checks to provide evidence that the relationship between expertise and arbitrage profit is not driven by a learning-by-doing mechanism. The estimation results are presented in Online Appendix Section A and will be discussed later.

²⁴ For completeness, we provide the trading patterns of other traders in Figure A.1.

²⁵ To further explore the relationship between these variables, we perform a principal component analysis (PCA), whose results are reported in Table A.1.

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 considerably larger amount of money on arbitrage than the former group of arbitrageurs. As the expected profit from trading is larger, one may expect that the level of effort is also higher. 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 differences in profit from arbitrage actions between the two groups of users is to estimate the following regression:

$$\text{Spread}_{i,j,p,t} = \beta_0 + \beta_1 \text{Trade Ability}_j + \beta_2 \text{USD}_{i,j,p,t} + \theta_p + \phi_t + \varepsilon_{i,j,p,t}, \quad (8)$$

where 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 the user level to account for redundant information across actions made by the same user.

The outcome, $\text{Spread}_{i,j,p,t}$, is the profit that a user j obtains by completing an arbitrage action i , using a dyad of currencies p , in percentage of the official exchange rate observed in the hour t . As described in Section 3.2, by construction the arbitrage action is profitable when the implied exchange rate is larger than the official exchange rate. The explanatory variable of interest, 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. (8) 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 because the latter is subject to high price volatility in time. To construct this variable, prices of the actions not in USD are converted to enable comparisons across currency markets. Most importantly, we include a set of currency pair (dyad) fixed effects, θ_p , which allow us to compare arbitrage actions operated using the same couple of currencies. We also introduce hourly time fixed effects, ϕ_t . As we have explained in Section 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 us to absorb any potential shocks that occurred on the market. In addition, by comparing arbitrage actions conducted in the very same hour, we likely capture contingent conditions of the market strictly related to risk, such as liquidity, volatility, and market depth, that otherwise would be difficult to capture given the “two-leg” (and “two-currency”) structure of the arbitrage actions.²⁶

In Table 6 we present our estimation results where trade ability is proxied by the dummy variable $D(\text{Currencies})$, 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 our 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) we add separately time and dyad fixed effects; in column (4) we add both fixed effects in the regression. Some observations are omitted 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 economically relevant: focusing on column (4), we 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.

In column (5) we test the robustness of this exercise and exclude users who perform just one arbitrage action. This group of users might comprise investors who engaged in arbitrage believing they would easily make risk-free profits, only to suddenly realize they could not, meaning they only acted once. Arguably, their inclusion might inflate the difference in profits we estimate between experts and non-experts. While this is partially true, as the estimated coefficient in column (5) is slightly smaller, the difference is still positive and statistically significant at the level of 1%. Further robustness checks are presented in Online Appendix. Table A.2 replicates column (4) for different selected parameters ΔQ and ΔT , whilst Table A.3 does the same for different values of ΔQ , holding fixed ΔT to 300 seconds. Moreover, Table A.4 repeats it when we identify a triangular arbitrage activity during a clear temporary opportunity (i.e., when the official exchange rate deviates above a given threshold). Reassuringly, all these checks demonstrate that the relationship between expertise and arbitrage profit does not depend on the way we identify the triangular arbitrage actions.²⁷

In addition, we show that this relationship does not depend on the specific proxy of users' expertise. With this purpose in mind, Table 7 replicates column (4) of Table 6 (which is reported in column (1) for easiness of comparison) using for expertise the following user-specific measures: the (logarithm of) the number of currencies used (column 2); the (logarithmic) number of arbitrage actions executed by the user (column 3); the dummy variable $D(\text{Metaorder})$, which is equal to 1 for all the actions conducted by users that executed at least one metaorder (column 4); the dummy variable $D(\text{Aggressive})$ that indicates whether the user executed at least one aggressive action (column 5); finally, in column 6 of Table 7, we use the scores of the first component obtained with the principal

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

²⁷ We note that when we set $\Delta Q = 0\%$, the sample drops by around half (column (3) of Table A.3). This might explain why the estimated coefficient, while still positive, loses statistical significance.

Table 6
Relationship between trade ability and profits.

Dep. var.:	Spread (with fees)				
	Baseline				$N_{act} > 1$
	(1)	(2)	(3)	(4)	
D(Currencies)	1.6180*** (0.1900)	1.5791*** (0.1943)	1.2421*** (0.1584)	1.2917*** (0.1659)	0.9832*** (0.2319)
Equiv. \$	3.4652** (1.7532)	2.6151* (1.5862)	0.6556 (1.2465)	0.1912 (1.1280)	0.4108 (1.1327)
Constant	-1.0420*** (0.1593)	-0.9985*** (0.1775)	-0.6141*** (0.1500)	-0.6506*** (0.1604)	-0.3558 (0.2253)
Time FE	N	N	hour	hour	hour
Dyad FE	N	Y	N	Y	Y
N	6594	6582	5307	5284	5176
R-squared	0.16	0.20	0.68	0.69	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. We 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. Column (5) is an additional specification excluding all the users that conducted only one arbitrage action ($N_{act} > 1$). All columns include an additional control for the amount of volume traded, expressed in USD (and divided by 10,000). We 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 7
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)	
PCI						0.2242*** (0.0466)
Equiv. \$	0.1912 (1.1280)	0.1525 (1.2621)	1.2060 (1.3191)	0.0733 (1.3384)	0.3375 (1.1244)	0.8175 (1.3500)
Constant	-0.6506*** (0.1604)	-1.0453 (0.7807)	-1.4603*** (0.4143)	0.3359** (0.1614)	0.6113*** (0.0424)	-1.0717*** (0.3593)
Time FE	hour	hour	hour	hour	hour	hour
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 6; (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) *PCI* is the score of each variable obtained by performing a PC analysis as explained in Table A.1. 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). We 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

component analysis. Overall, these additional results are consistent with a positive relationship between expertise and arbitrage profits.²⁸

²⁸ Furthermore, in Online Appendix we show that these additional estimation results are, to a large extent, not sensitive to varying the boundaries for the identification of the triangular arbitrage action (i.e., ΔQ and ΔT). See Tables A.5 and A.6. Similar results are found if we consider only arbitrage actions executed when the exchange rate exceeds a certain threshold.

Finally, we show that this relationship is not likely to be driven by a “learning-by-doing” process. For example, excluding from the sample the users who start investing using a new dyad of currencies only after a relatively long period (of fourteen days) yields estimations that are consistent with those presented in Table 7, regardless of the measures of expertise we use (see Online Appendix Table A.12). We obtain similar conclusions if we further exclude the users who were active on Mt. Gox before September 2011 (see Table A.14).

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. Why is this the case? In this section we show that the differences in profits stem from a better ability among the former to respond quickly to fluctuations, which makes arbitrage more (or less) profitable. Indeed, a typical profitable situation in financial markets arises when unexpected deviations occur in fundamental values. Due to structural frictions, adjustments across markets are not automatic, giving rise to opportunities to conduct arbitrage operations. We exploit this fact in our analysis and reconstruct, from the hourly evolution of the official exchange rate in a market, the unsigned percent variation in the exchange rate with respect to the previous hour (see Table 1 for the formal definition). Our 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 , changes 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, rather, requires expertise and/or the execution of automated orders. Hence, when variation in the exchange rate is more prominent than in typical times it is likely that expert users take advantage of this to make profits. Second, as the users who trade are small (as we will discuss later, there are likely no institutional investors operating on Mt. Gox), their actions are unlikely to affect such deviations. It is therefore reasonable to assume that users are exchange rate *takers* and that deviations in the official exchange rate are exogenous.

We employ this variable $|\Delta R|_{p,t}$ on the right-hand side of our 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:

$$\text{Spread}_{i,j,p,t} = \beta_1(\text{Trade Ability}_j \times |\Delta R|_{p,t}) + \beta_2|\Delta R|_{p,t} + \beta_3 \text{USD}_{i,j,p,t} + \alpha_j + \theta_p + \phi_t + \epsilon_{i,j,p,t}. \tag{9}$$

As one can see, our main variable of interest in Eq. (9) is now time variant. This is important as it allows us to employ a set of user fixed effects, α_j , which permit us 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 our measure of trade ability. The inclusion of α_j also implies that our 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% increase in (the absolute value of the) rate of change of the official exchange rate. β_2 captures the effect of a 1% increase in (the absolute value of 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 hour 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 the user-level as above.

Table 8 reports the estimates of the main coefficients of interest using different measures of trade ability. In the first two columns, we use the variable $D(\text{Currencies})$. We then repeat the analyses by using the alternative proxies of expertise that have been used above. Overall, we find that an increase in the (absolute value of 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, we note that 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 — indeed, as we showed in Table 3, many of the users active in a single market executed just one action. This leads to the exclusion of a significant number of observations from this group, making it more difficult to obtain stable and statistically significant results. Finally, the result for column (10) — relative to the aggressiveness of the actions — goes against our expectations, but the coefficient is statistically imprecise. Overall, these findings indicate that sophisticated investors are more able to take into account and exploit in their favor quick changes in the official exchange rate than are non-expert users, and that this ability leads to higher profits. For example, looking at the effect reported in column (12) of Table 8, we estimate that an arbitrageur with a trade ability score that is a standard deviation above the mean obtained a profit of 1.347% following a 1% increase in the rate of change of the official exchange rate (i.e., 0.476×2.83); note that this premium accounts for more than a standard deviation of the dependent variable. Our interpretation, indeed, is that the expert arbitrageurs are more able with respect to the others to react to price deviations, and thus their activity is also more profitable.

Varying the way in which we identify a triangular arbitrage action yields estimate results that are qualitatively comparable. When we use a more conservative algorithm we obtain similar, if not more statistically robust, estimates (for example, Table A.7 replicates our analysis setting $\Delta Q = 1\%$ and $\Delta T = 30$ s, while Table A.8 does it with $\Delta Q = 0\%$ and $\Delta T = 300$ s).²⁹ Conversely (and in line with

²⁹ We note that when we use $D(\text{Currencies})$ as the proxy of expertise and set a boundary for the volume of the triangular arbitrage that is smaller or equal to 1% (i.e., $\Delta Q \leq 1\%$) the sample shrinks by more than one quarter and the estimated coefficient $\hat{\beta}_2$ drops. This is potentially explained by the fact that we might have few traders that execute arbitrage just using a couple of fiat currencies (in a given hour) when we identify these actions perfectly (i.e., in a tight neighborhood of ΔQ). We also point out that when we use continuous variables to measure expertise, like the principal component, we still have sufficient variation in trade sophistication to estimate $\hat{\beta}_2$.

Table 8
Responsiveness to official rate variations.

Dep. var.:	Spread (with fees)												
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	
$ \Delta R \times D(\text{Currencies})$	6.905*** (1.139)	1.442 (5.314)											
$ \Delta R \times \text{Log}(\text{Currencies})$			3.213** (1.528)	1.703** (0.723)									
$ \Delta R \times \text{Log}(\text{Actions})$					1.427*** (0.202)	0.150 (0.409)							
$ \Delta R \times D(\text{Metaorder})$							3.693** (1.600)	2.228*** (0.821)					
$ \Delta R \times D(\text{Aggressive})$									-6.786*** (1.595)	3.593 (3.461)			
$ \Delta R \times \text{PCI}$											1.048*** (0.159)	0.476** (0.211)	
$ \Delta R $	-5.441*** (1.196)	-0.742 (5.279)	-4.744** (2.125)	-2.426 (1.590)	-7.060*** (1.320)	-0.258 (2.423)	-2.341* (1.241)	-1.045 (0.854)	0.168 (0.526)	0.617 (0.956)	-5.791*** (1.187)	-2.810* (1.546)	
Equiv. \$	1.150 (1.700)	-0.732 (0.833)	0.732 (1.762)	-0.861 (0.878)	2.128 (1.815)	-0.751 (0.827)	0.766 (1.647)	-0.804 (0.871)	0.824 (1.896)	-0.741 (0.838)	2.225 (1.800)	-0.812 (0.866)	
User FE	N	Y	N	Y	N	Y	N	Y	N	Y	N	Y	
Time FE	N	hour	N	hour	N	hour	N	hour	N	hour	N	hour	
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	

Notes: the Table describes the responsiveness to variations of the official exchange rate for the main proxies of trade ability, and their effect on profits. It reports OLS estimates for 12 different specifications, each including an interaction term between the official rate variation and a proxy of trade ability: D(Currencies) in (1-2), Log(Currencies) in (3-4), Log(Actions) in (5-6), D(Metaorder) in (7-8), D(Aggressive) in (9-10), PCI in (11-12). $|\Delta R|$ is unsigned, i.e. it is in absolute values. The first column of each alternative proxy is without fixed effects, while the second includes time, dyad and user fixed effects. All columns include a control for the amount of volume traded, expressed in USD (and divided by 10,000), and for the official rate variation. We 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

our expectations), estimates become more imprecise when we use a more inclusive algorithm (i.e., with $\Delta Q = 20\%$ and $\Delta T = 600s$) that likely adds noise in the identification of arbitrage trades (Table A.9). Reassuringly, our main conclusions also hold when we further restrict the identification along other relevant dimensions. For example, by restricting a triangular arbitrage action to those executed in an hour in which the official exchange rate fluctuates prominently (i.e., $|\Delta R| > r$, with $r = \{0.01\%, 0.05\%, 0.1\%\}$), we still obtain a positive and statistically significant premium for the expert users that (as expected) is larger the more we zoom in on more profitable timeframes (Table A.10).³⁰

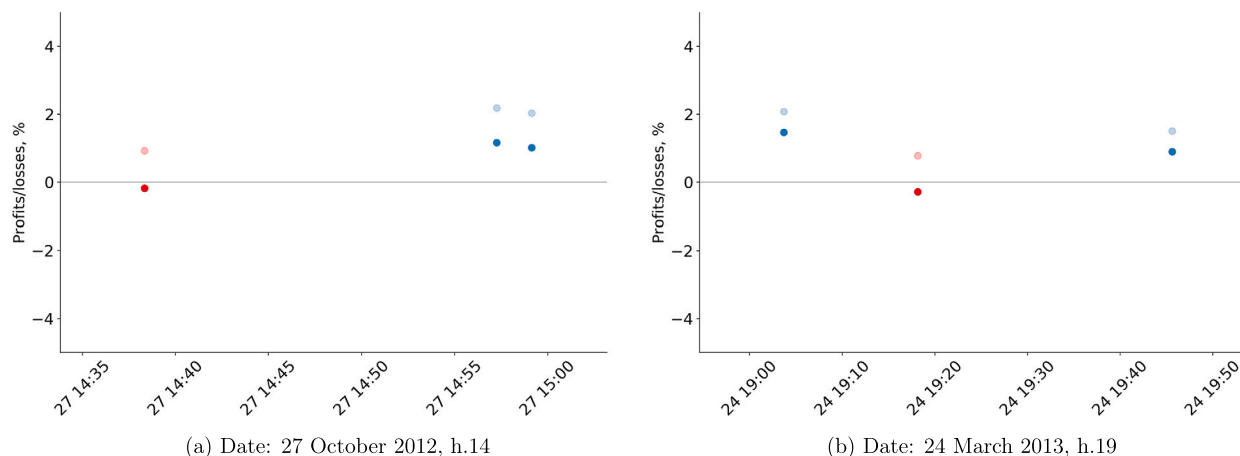
Finally, we turn to the possible alternative story of users learning how to perform arbitrage by trading. To be sure that this channel does not explain our findings, in Online Appendix we present a battery of robustness checks. We first replicate Table 8 with user fixed effects by removing from the sample the users who search for new arbitrage opportunities on new currency markets after a relatively large period (Table A.13); we then remove the users who operated on Mt. Gox before September 2011 (Table A.15). In both exercises, we find estimates that are comparable to our baseline analysis. Even more importantly, we show that our results hold even when relaxing the assumption that expertise is fixed in a user. In Table A.16, for example, we include user-by-month fixed effects in place of user fixed effects. This means that we identify the effect of responsiveness to arbitrage opportunity on arbitrage profits by exploiting variation across the trades executed by a user in a given month. To do this, we assume that *within a given month* the user’s expertise is fixed. Reassuringly, this exercise also hints at a relationship between expertise and arbitrage profit that is not explained by alternative mechanisms, like the “learning-by-doing” process.

6. Why is non-expert users’ arbitrage unprofitable?

Our analysis indicates that expert users obtain on average a positive profit from arbitraging. Looking closer at the regression tables, we also note that, throughout all the specifications we test, the constant term — β_0 — in the regression Eq. (8) is either negatively estimated or not statistically significant. This indicates that non-expert users’ arbitrage is on average unprofitable. In this section, we try to shed light on the reasons for this apparently puzzling finding.

Table 8 suggests that, unlike the expert users, non-sophisticated arbitrageurs have a limited ability to exploit quick and temporary deviations in the exchange rates. What is more, they seem to make mistakes when prices fluctuate considerably (we actually estimate a negative β_2 , not statistically significant in some columns). This is likely because they do not incorporate the relevant information in their investment strategies (neither using APIs nor automated trading algorithms) and make mistakes in choosing *when* conducting the arbitrage actions: to reiterate, our evidence underlines the importance of the timing of execution at the micro scale, which in

³⁰ Our results do not change if we further restrict arbitrage actions to those with $\Delta Q = 0\%$ in addition to $|\Delta R| > r$. See Table A.11, which replicates the analysis conducted in Table A.10.



Notes: the plots show the arbitrage activity in two different time windows. Each action is reported twice, once including and once excluding the transaction costs (the former is slightly 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, and once the fees are included their actions yield losses.

Fig. 7. The ‘monetary illusion’ effect.

turn determines a crucial difference between a profitable and a non-profitable action. Further inspections provide limited evidence that losses are moderately higher when arbitraging with a larger amount of USD dollars; instead, we do not find evidence that a retarded action (i.e., a large ΔT) or a large mismatch in the trade’s volume (i.e., a large ΔQ) explain why this group of users obtain losses from arbitraging (see Online Appendix Table A.17).

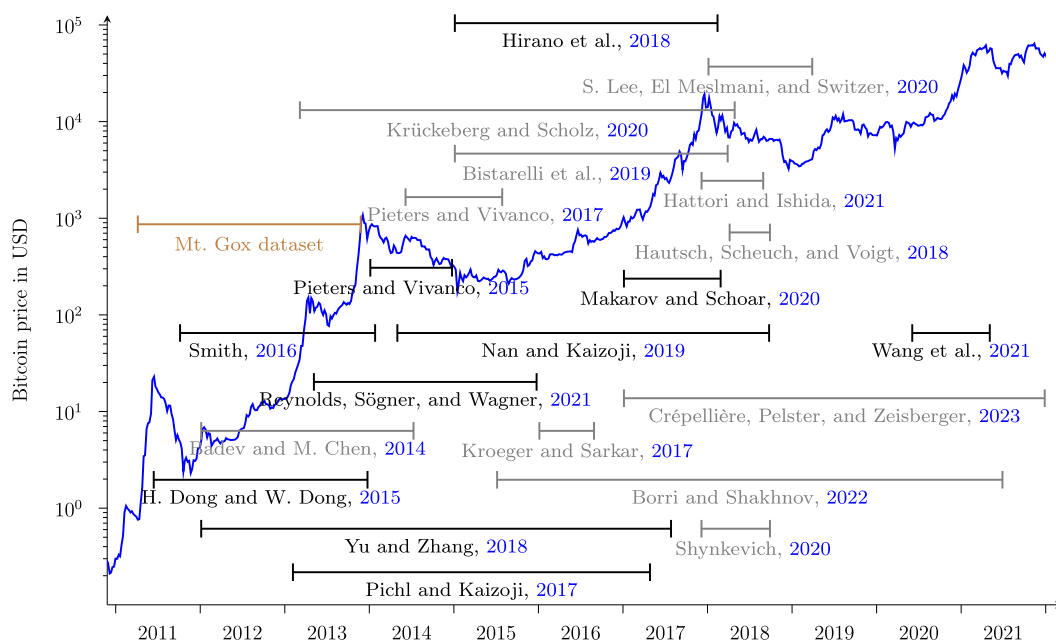
Our data also suggest a potential explanation that is based on the transaction costs users have to pay to the exchange. In Online Appendix Table B.4, for example, we repeat this analysis by computing the spreads with alternative measures of the transaction costs, with the results indicating that the constant term is non-negative only when the transaction costs are not taken into account. It is therefore possible that the non-sophisticated users do not account correctly for the costs of conducting the arbitrage strategy in a mechanism akin to the one originating the monetary illusion phenomenon (Shafir et al., 1997) and thus incur unprofitable activity. Fig. 7 provides an example for two different illustrative time windows (27 October 2012 and 24 March 2013), which is consistent with this interpretation: 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 engaged in arbitrage in a single market, while the blue dots denote the actions executed by expert investors active in multiple 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, the actions are in general less profitable and even unprofitable once the transaction costs are included, both when they are similar across users (Panel 7a) and when they vary across them (Panel 7b).

7. Interpretation in temporal and market context

Before presenting our conclusions, we discuss how our results might be interpreted if applied to other temporal and market contexts. To this end, we begin by noting that the body of previous literature that applies financial econometrics to time series data from cryptocurrency exchanges is vast and not easy to navigate.³¹ Most studies use relatively short and often non-overlapping samples, which makes it difficult to derive general conclusions regarding arbitrage conducted in cryptocurrency markets. This is stylized in Fig. 8 where we report a selection of prior works on arbitrage using bitcoins (in gray bands) and on triangular arbitrage in particular (in black bands). The temporal sample these works investigate is represented by the x-axis. Focusing on the latter group, one can see that some studies (e.g., Smith, 2016; Dong and Dong, 2015) have analyzed the rising period of Bitcoins (2011–2013); the majority of works we find in the literature analyzed arbitrage during the years 2014–2018 (e.g. Hirano et al., 2018; Pieters and Vivanco, 2015; Makarov and Schoar, 2020; Nan and Kaizoji, 2019; Reynolds et al., 2021; Pichl and Kaizoji, 2017), while the recent years have arguably been understudied.

In the y-axis of Fig. 8, we draw the Bitcoin price in USD on a logarithmic scale to better account for differences across epochs. By doing this, it becomes easy to spot even graphically multiple structural breaks in the time series. The value of the Bitcoins has increased rapidly over the USD — by 5 orders of magnitude in just 10 years. This has changed the structure of the cryptocurrency markets, including the market size and the composition of traders as well as the probability of getting temporary arbitrage opportunities. The available studies trace some general trends that help us to understand in which directions these markets have changed. On the one hand, trade volumes have significantly increased over the years and so too have the number of websites and cryptocurrency

³¹ E.g., Glaser et al. (2014); Garcia et al. (2014); Brandvold et al. (2015); Yermack (2015); Cheung et al. (2015); Ciaian et al. (2016); Athey et al. (2016); Bouri et al. (2017); Katsiampa (2017); Wheatley et al. (2019); Dyhrberg et al. (2018).



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.

Fig. 8. Related work on arbitrage in temporal and market context (links for bibliography: Badev and Chen, 2014; Bistarelli et al., 2019; Hattori and Ishida, 2021; Hautsch et al., 2018; Kroeger and Sarkar, 2017; Krückeberg and Scholz, 2020; Lee et al., 2020; Pieters and Vivanco, 2017; Shynkevich, 2020).

platforms (e.g., Borri and Shakhnov, 2022). On the other hand, arbitrage opportunities have declined (Crépeillère et al., 2023), despite volatility in the Bitcoin price not falling (e.g., Bourghelle et al., 2022).

Our findings are drawn from the analysis of a nascent cryptocurrency market (the orange band in Fig. 8 draws our sample period). In light of these changes, one may wonder how easily generalizable our conclusions are and how they might apply to a more mature market context like the current cryptocurrency market. We can see two directions of bias in our estimates (if compared to a current, changed world). For one thing, given that trade volumes and exchanges have increased, one should expect to observe many more investors trading Bitcoins, potentially, without any specific financial skills. One should therefore expect the share of expert users to be lower than the share we observe in our context. Moreover, since arbitrage opportunities have declined over recent years, one should expect expert investors to have become more sophisticated and responsive. In other words, one should expect there to be fewer, yet more sophisticated, expert arbitrageurs in today's market. On top of that, we highlight that during the earlier years of the cryptocurrency market, traders investing in bitcoin were self-selected traders with the minimum level of sophistication required to figure out how cryptocurrencies worked. This is at odds with the current context, in which we observe a population of Bitcoin investors who possess a good general knowledge of what cryptocurrencies are and how they work. This means that one should expect non-expert users today to be less sophisticated than in the context we study. These two facts suggest that our estimates might be downward biased.

For another one, the rapid increase in the number of websites and cryptocurrency platforms now provides more information to traders. Investors' searching costs in detecting arbitrage opportunities are therefore expected to be lower today than in our context — a change that could make expertise today a less important feature in conducting arbitrage (which suggests that our estimates might be upward biased). In different contexts, one effect may dominate the other. Understanding which one of the two prevails can help us to understand the magnitude of the impact of expertise on arbitrage profits.

8. Concluding remarks

In this paper we use trader-specific information from the leaked dataset of a Bitcoin trading platform, Mt. Gox, to identify the triangular arbitrage activity and to investigate whether the arbitrageurs are many small traders or a few sophisticated investors. The conventional economic interpretation of theoretical arbitrage 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 sufficiently small exposure on the market. Anecdotal evidence indicates that arbitrageurs are few, sophisticated, and specialized traders. Our analysis provides empirical evidence to support this statement. We find that sophisticated arbitrageurs make systematic profits from these triangular actions and that the distribution of their profits first-order stochastically dominates that achieved by the non-expert arbitrageurs. They do so by reacting more effectively to exogenous shocks, such as temporary movements in the official exchange rate. We show that our conclusions are not sensitive to a number of robustness checks, such as alternative identification of a triangular arbitrage action or alternative samples.

While our findings are drawn from the analysis of a singular cryptocurrency market, Mt. Gox, there are elements that suggest that our results are also valid in traditional markets or similar contexts. For instance, the arbitrage activity in Mt. Gox was likely conducted by individual traders and not by specialized professionals who operate on someone else's funds (as in Shleifer and Vishny (1997)). Despite this, our data document considerable variation in the level of expertise and a premium (in terms of profits) in favor of the most sophisticated users. Moreover, recent findings based on the Ethereum ecosystem confirm that expert users' actions are more profitable because they are conducted using sophisticated methods like private smart contracts that require deep knowledge of the working ecosystem (Wang et al., 2022). In this paper, we discuss several contextual factors that might help the reader to interpret how our conclusions can apply to a more mature cryptocurrency market.

Some limitations to our work stand out. Firstly, our data do not cover the user-specific characteristics of the arbitrageurs, which prevents us from conducting a more general analysis linking demographic features to arbitrage. Moreover, we do not know whether the users we observe participate in other trading activities, including on other markets. Triangular arbitrage aligns prices in one market, whereas an essential function of arbitrage is its function of "information carrier" across markets. While the evidence might be stronger for triangular arbitrage within the same market, our results may provide a broader picture if complemented by studies investigating the behavior of arbitrageurs conducting other arbitrage strategies that are either more complex or executed across multiple markets.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We would like to thank two anonymous reviewers, Andrea Canidio, Simone Giansante, JT Hamrick, Roberto Renò, Antonio Scala, Marie Vasek, Jiahua Xu, as well as the participants of the University of Innsbruck's Security & Privacy research seminar and the participants of the Workshop on the Economics of Information Security 2019 (WEIS) for their comments on earlier versions of this work (Saggese and Böhme, 2019). Angelo Facchini thankfully acknowledges financial support by the European Union Horizon 2020 Program under the scheme INFRAIA-01-2018-2019 Integrating Activities for Advanced Communities, Grant Agreement n. 871042, SoBigData++: European Integrated Infrastructure for Social Mining and Big Data Analytics.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2023.07.025>.

References

- Aliber, Robert Z., 1973. The interest rate parity theorem: a reinterpretation. *J. Polit. Econ.* 81 (6), 1451–1459.
- Aloosh, Arash, Li, Jiasun, 2019. Direct evidence of bitcoin wash trading. Available at SSRN 3362153.
- Athey, Susan, et al., 2016. Bitcoin pricing, adoption, and usage: theory and evidence. Available at SSRN 2826674.
- Badev, Anton, Chen, Matthew, 2014. Bitcoin: technical background and data analysis. In: *Finance and Economics Discussion Series*. Federal Reserve Board, pp. 2014–2104.
- Barberis, Nicholas, Thaler, Richard H., 2003. A survey of behavioral finance. In: *Handbook of the Economics of Finance* 1, pp. 1053–1128.
- Bistarelli, Stefano, et al., 2019. Model-based arbitrage in multi-exchange models for Bitcoin price dynamics. *Digit. Finance* 1 (1–4), 23–46.
- Black, Fischer, Scholes, Myron, 1973. The pricing of options and corporate liabilities. *J. Polit. Econ.* 81, 637–654.
- Böhme, Rainer, et al., 2015. Bitcoin: economics, technology, and governance. *J. Econ. Perspect.* 29 (2), 213–238.
- Borri, Nicola, Shakhnov, Kirill, 2022. The cross-section of cryptocurrency returns. *Rev. Asset Pricing Stud.* 12 (3), 667–705.
- Borsboom, Charlotte, Zeisberger, Stefan, 2020. What makes an investment risky? An analysis of price path characteristics. *J. Econ. Behav. Organ.* 169, 92–125.
- Bourghelle, David, Jawadi, Fredj, Rozin, Philippe, 2022. Do collective emotions drive bitcoin volatility? A triple regime-switching vector approach. *J. Econ. Behav. Organ.* 196, 294–306.
- Bouri, Elie, Azzi, Georges, Dyrhberg, Anne, 2017. On the return-volatility relationship in the Bitcoin market around the price crash of 2013. *Economics* 11 (1).
- Brandvold, Morten, et al., 2015. Price discovery on Bitcoin exchanges. *J. Int. Financ. Mark. Inst. Money* 36, 18–35.
- Budish, Eric, Cramton, Peter, Shim, John, 2015. The high-frequency trading arms race: frequent batch auctions as a market design response. *Q. J. Econ.* 130 (4), 1547–1621.
- 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, Roca, Eduardo, Su, Jen-Je, 2015. Crypto-currency Bubbles: an application of the Phillips-Shi-Yu (2013) Methodology on Mt. Gox Bitcoin Prices. *Appl. Econ.* 47 (23), 2348–2358.
- Ciaian, Pavel, Rajcaniova, Miroslava, Kancs, d'Artis, 2016. The economics of BitCoin price formation. *Appl. Econ.* 48 (19), 1799–1815.
- Crépellière, Tommy, Pelster, Matthias, Zeisberger, Stefan, 2023. Arbitrage in the market for cryptocurrencies. *J. Financ. Mark.*, 100817.
- De Bortoli, Daiane, et al., 2019. Personality traits and investor profile analysis: a behavioral finance study. *PLoS ONE* 14 (3), e0214062.
- De Long, Bradford J., et al., 1990. Noise trader risk in financial markets. *J. Polit. Econ.* 98 (4), 703–738.
- Dimitri, Nicola, 2019. Transaction fees, block size limit, and auctions in Bitcoin. *Ledger* 4.

- Dong, Huijian, Dong, Weigu, 2015. Bitcoin: exchange rate parity, risk premium, and arbitrage stickiness. *J. Econ. Manag. Trade*, 105–113.
- Donier, Jonathan, Bonart, Julius, 2015. A million metaorder analysis of market impact on the Bitcoin. *Mark. Microstruct. Liq.* 1 (02).
- Dyhrberg, Anne, Foley, Sean, Svec, Jiri, 2018. How investible is Bitcoin? Analyzing the liquidity and transaction costs of Bitcoin markets. *Econ. Lett.* 171, 140–143.
- Easley, David, O'Hara, Maureen, Basu, Soumya, 2019. From mining to markets: the evolution of bitcoin transaction fees. *J. Financ. Econ.* 134 (1), 91–109.
- Fama, Eugene, 1970. Efficient capital markets: a review of theory and empirical work. *J. Finance* 25 (2), 383–417.
- Feder, Amir, et al., 2018. The impact of DDoS and other security shocks on Bitcoin currency exchanges: evidence from Mt. Gox. *J. Cybersecurity* 3 (2), 137–144.
- Froot, Kenneth A., Dabora, Emil M., 1999. How are stock prices affected by the location of trade? *J. Financ. Econ.* 53 (2), 189–216.
- Gandal, Neil, et al., 2018. Price manipulation in the Bitcoin ecosystem. *J. Monet. Econ.* 95, 86–96.
- Garcia, David, et al., 2014. The digital traces of bubbles: feedback cycles between socio-economic signals in the Bitcoin economy. *J. R. Soc. Interface* 11 (99).
- Glaser, Florian, et al., 2014. Bitcoin-asset or currency? Revealing users' hidden intentions. In: *ECIS European Conference on Information Systems*.
- Griffin, John M., Shams, Amin, 2020. Is Bitcoin really untethered? *J. Finance* 75 (4), 1913–1964.
- Gromb, Denis, Vayanos, Dimitri, 2010. Limits of arbitrage. *Annu. Rev. Financ. Econ.* 2 (1), 251–275.
- Halaburda, Hanna, et al., 2022. The microeconomics of cryptocurrencies. *J. Econ. Lit.* 60 (3), 971–1013.
- Harris, Lawrence, Gurel, Eitan, 1986. Price and volume effects associated with changes in the S&P 500 list: new evidence for the existence of price pressures. *J. Finance* 41 (4), 815–829.
- Harrison, J. Michael, Kreps, David M., 1979. Martingales and arbitrage in multiperiod security markets. *J. Econ. Theory* 20, 381–408.
- Haryanto, Steven, Subroto, Athor, Ulpah, Maria, 2019. Disposition effect and herding behavior in the cryptocurrency market. *J. Ind. Bus. Econ.*, 1–18.
- Hasbrouck, Joel, Saar, Gideon, 2013. Low-latency trading. *J. Financ. Mark.* 16 (4), 646–679.
- Hasso, Tim, Pelster, Matthias, Breitmayer, Bastian, 2019. Who trades cryptocurrencies, how do they trade it, and how do they perform? Evidence from brokerage accounts. *J. Behav. Exp. Finance* 23, 64–74.
- Hattori, Takahiro, Ishida, Ryo, 2021. The relationship between arbitrage in futures and spot markets and Bitcoin price movements: evidence from the Bitcoin markets. *J. Futures Mark.* 41 (1), 105–114.
- Hautsch, Nikolaus, Scheuch, Christoph, Voigt, Stefan, 2018. Limits to arbitrage in markets with stochastic settlement latency. In: *CFS Working Paper*.
- Hirano, Yuki, et al., 2018. Analysis of Bitcoin market efficiency by using machine learning. In: *CBU International Conference Proceedings*, vol. 6, pp. 175–180.
- Katsiampa, Paraskevi, 2017. Volatility estimation for Bitcoin: a comparison of GARCH models. *Econ. Lett.* 158, 3–6.
- Kirilenko, Andrei A., Lamacie, Gui, 2015. Latency and Asset Prices. Available at SSRN 2546567.
- Kourtidis, Dimitrios, Šević, Željko, Chatzoglou, Prodromos, 2011. Investors' trading activity: a behavioural perspective and empirical results. *J. Socio-Econ.* 40 (5), 548–557.
- Kroeger, Alexander, Sarkar, Asani, 2017. Law of one bitcoin price? Tech. Rep. Federal Reserve Bank of Philadelphia.
- Krückeberg, Sinan, Scholz, Peter, 2020. Decentralized efficiency? Arbitrage in Bitcoin markets. *Financ. Anal. J.* 76 (3), 135–152.
- Lamont, Owen A., Thaler, Richard H., 2003. Anomalies: the law of one price in financial markets. *J. Econ. Perspect.* 17 (4), 191–202.
- Lee, Charles M.C., Radhakrishna, Balkrishna, 2000. Inferring investor behavior: evidence from TORQ data. *J. Financ. Mark.* 3 (2), 83–111.
- Lee, Charles M.C., Ready, Mark J., 1991. Inferring trade direction from intraday data. *J. Finance* 46 (2), 733–746.
- Lee, Seunggho, El Meslmani, Nabil, Switzer, Lorne N., 2020. Pricing efficiency and arbitrage in the Bitcoin spot and futures markets. *Res. Int. Bus. Finance* 53.
- Lintner, John, 1965. Security prices, risk and maximal gains from diversification. *J. Finance* 20 (4), 587–615.
- von Luckner, C.G., Reinhardt, C.M., Rogoff, K., 2023. Decrypting new age international capital flows. *J. Monet. Econ.*
- Makarov, Igor, Schoar, Antoinette, 2020. Trading and arbitrage in cryptocurrency markets. *J. Financ. Econ.* 135 (2), 293–319.
- Makridis, Christos A., Fröwis, Michael, Sridhar, Kiran, Böhme, Rainer, 2023. The rise of decentralized cryptocurrency exchanges: Evaluating the role of airdrops and governance tokens. *J. Corp. Finance* 79 (4).
- Malkiel, Burton G., 2003. The efficient market hypothesis and its critics. *J. Econ. Perspect.* 17 (1), 59–82.
- Merton, Robert C., 1973. Theory of rational option pricing. *Bell J. Econ. Manag. Sci.*, 141–183.
- Moore, Tyler, Christin, Nicolas, Szurdi, Janos, 2018. Revisiting the risks of Bitcoin currency exchange closure. *ACM Trans. Internet Technol.* 18 (4), 50.
- Möser, Malte, Böhme, Rainer, 2015. Trends, tips, tolls: a longitudinal study of Bitcoin transaction fees. In: *International Conference on Financial Cryptography and Data Security*. Springer, pp. 19–33.
- Mossin, Jan, 1966. Equilibrium in a capital asset market. *Econometrica* 34 (4), 768–783.
- Nan, Zheng, Kaizoji, Taisei, 2019. Bitcoin-based triangular arbitrage with the Euro/US dollar as a foreign futures hedge: modeling with a bivariate GARCH model. *Q. Finance Econ.* 3 (2), 347–365.
- Petrov, Andreas, Schuffla, David, 2013. Arbitrage am Bitcoin-Markt. In: Böhme, Rainer, Breuker, Dominic (Eds.), *1st Münster Bitcoin Conference*. University of Münster.
- Pichl, Lukáš, Kaizoji, Taisei, 2017. Volatility analysis of Bitcoin price time series. *Q. Finance Econ.* 1 (4), 474–485.
- Pieters, Gina, Vivanco, Sofia, 2015. Bitcoin arbitrage and unofficial exchange rates. Tech. Rep. Trinity University.
- Pieters, Gina, Vivanco, Sofia, 2017. Financial regulations and price inconsistencies across Bitcoin markets. *Inf. Econ. Policy* 39, 1–14.
- Reynolds, Julia, Sögner, Leopold, Wagner, Martin, 2021. Deviations from triangular arbitrage parity in foreign exchange and bitcoin markets. *Cent. Eur. J. Econ. Model. Econom.*, 105–146.
- Roll, Richard, Ross, Stephen A., 1980. An empirical investigation of the arbitrage pricing theory. *J. Finance* 35 (5), 1073–1103.
- Ross, Stephen A., 1976. The arbitrage theory of capital asset pricing. *J. Econ. Theory* 13 (3), 341–360.
- Saggese, Pietro, Böhme, Rainer, 2019. Identifying the arbitrageurs on Mt. Gox: First insights from the leaked dataset. In: *The 18th Workshop on the Economics of Information Security (WEIS)*.
- Scaillet, Olivier, Treccani, Adrien, Trevisan, Christopher, 2020. High-frequency jump analysis of the bitcoin market. *J. Financ. Econ.* 18 (2), 209–232.
- Shafir, Eldar, Diamond, Peter, Tversky, Amos, 1997. Money illusion. *Q. J. Econ.* 112 (2), 341–374.
- Sharpe, William F., 1964. Capital asset prices — a theory of market equilibrium under conditions of risk. *J. Finance* 19 (3), 425–442.
- Shleifer, Andrei, Vishny, Robert W., 1997. The limits of arbitrage. *J. Finance* 52 (1), 35–54.
- Shynkevich, Andrei, 2020. Bitcoin arbitrage. *Finance Res. Lett.*, 293–319.
- Smith, Jacob, 2016. An Analysis of Bitcoin Exchange Rates. Available at SSRN 2493797.
- Stoll, Hans R., 2000. Presidential address: friction. *J. Finance* 55 (4), 1479–1514.
- Wang, Y., Chen, Y., Wu, H., Zhou, L., Deng, S., Wattenhofer, R., 2022. Cyclic arbitrage in decentralized exchanges. In: *Companion Proceedings of the Web Conference 2022*, pp. 12–19.
- Westphal, Rebecca, Sornette, Didier, 2020. Market impact and performance of arbitrageurs of financial bubbles in an agent-based model. *J. Econ. Behav. Organ.* 171, 1–23.
- Wheatley, Spencer, et al., 2019. Are Bitcoin bubbles predictable? Combining a generalized Metcalfe's law and the log-periodic power law singularity model. *R. Soc. Open Sci.* 6 (6).
- Yermack, David, 2015. Is Bitcoin a real currency? An economic appraisal. In: *Handbook of Digital Currency*. Elsevier, pp. 31–43.
- Yu, Gloria, Zhang, Jinyuan, 2018. A revisit to capital control policies when Bitcoin is in town. In: *31st Australasian Finance and Banking Conference*.