

How Cryptocurrency Exchange Interruptions Create Arbitrage Opportunities

Andrew Morin
Tandy School of Computer Science
The University of Tulsa
Tulsa, OK, USA
amorin@utulsa.edu

Tyler Moore
School of Cyber Studies
The University of Tulsa
Tulsa, OK, USA
tyler-moore@utulsa.edu

Abstract—Centralized cryptocurrency exchanges offer users a more convenient platform to trade their digital assets at the cost of reduced control. As a result, when these exchanges suffer interruptions users struggle to access their funds or modify their orders. We investigate 41 events at the popular exchange Bitfinex, and measure the impact these events have on trades, volume, and pricing. We find that the volume to trade ratio increases during events, as fewer traders are moving large amounts of bitcoin. We also find that these interruptions often occur at the same time as arbitrage opportunities, with substantial profit opportunities.

1. Introduction

Despite decentralization being one of the core tenets of cryptocurrencies, the majority of users are eager to avoid the hassle. Rather than trading cryptocurrencies at a decentralized exchange, which preserves much of blockchain’s inherent traits, most trading occurs at centralized exchanges. Users of such exchanges enjoy near-instant trades, low fees, and high liquidity, all by simply abandoning the blockchain. The exchange centrally manages all transactions, and only incorporates the blockchain when interacting with external agents. Unfortunately, this centralized custodian-based architecture produces as many problems as it solves. Storing vast amounts of user funds in a handful of accounts creates a lucrative target for cybercriminals [1]–[3], and trusting exchanges to be responsible custodians in the absence of meaningful consumer protections or regulation leads to exchange failures such as QuadrigaCX or FTX.

An additional consequence of this departure from the blockchain, is the inability for users to access their funds during an exchange interruption. Because these exchanges are, at their core, a cryptocurrency wallet and database connected to a website, they are susceptible to the challenges of any other website: outages, DDoS attacks, third-party delays, etc. What’s more, because cryptocurrency exchanges exist largely in a regulatory vacuum, they avoid the rigorous obligations of their traditional finance counterparts, which are designed to forestall these problems. While a DDoS attack suffered by New Zealand’s Exchange and an outage of NASDAQ both prompted regulatory investigations [4], [5], the frequent interruptions of cryptocurrency exchanges rarely result in more than a Twitter post.

When an exchange suffers an interruption, the ability for users to create or modify trades is hindered, yet rarely ceases entirely. For example, the website may become inaccessible to human users, yet the back-end servers are still processing automated trades placed in advance. Whatever the cause, this results in a fragmented market with the majority of buy and sell orders at the affected exchange being essentially frozen in time as the other exchanges continue normal trading activity. This leads to a price difference between markets, and an arbitrage opportunity for those lucky few who are still able to trade on the interrupted exchange.

Arbitrage is a popular trading strategy which takes advantage of small price differences of the same good in different markets. By purchasing a cryptocurrency at a discounted price at one exchange, and transferring them to a second exchange buying at a premium, arbitrageurs can make a minor profit on the price difference. Because the difference in price is often small, a tangible profit requires a large volume of trades. This profit strategy is another well researched field as it pertains to cryptocurrencies [6], [7]. Makarov and Schoar find that cryptocurrency markets do exhibit periods of large arbitrage opportunities, however they are more common across geographic regions where capital controls prevent the movement of arbitrage profits. Two major factors limiting arbitrage opportunities outlined by Makarov and Schoar, are related to the long duration required to move cryptocurrency over the blockchain¹, and the restrictions imposed on trading between government-issued currencies.

The arbitrage opportunities created during exchange interruptions are confined to a single centralized exchange, avoiding all blockchain activity and capital controls. Additionally, Krückeberg et al. specifically identify Bitfinex as a lucrative target for buy-side arbitrage opportunities due to frequent discounted pricing [6]. In this paper, we investigate the arbitrage opportunities created during these exchange interruptions.

The rest of the paper will continue as follows: in Section 2 we review related work. In Section 3 we discuss the data collected, and the steps taken in our analysis of interruption events. In Section 4 we display our findings, and in Section 5 we provide a brief review.

1. Bitcoin transfers are generally considered valid after several block confirmations, which on average take 5-10 minutes each.

2. Related Work

Cryptocurrency trade strategies and manipulations is a heavily researched topic. Eigelshoven et al. [8] and Twomey et al. [9] both provide extensive reviews of existing literature on this topic. While some papers highlight unique and perplexing behaviors of cryptocurrencies, such as bizarre price and volume behaviors [10], or potential scams capable of sinking the entire market, [11], most draw parallels to long understood traditional market behaviors. This is due in no small part to cryptocurrency participants tendency to gravitate towards traditional finance market designs. Derivative assets, centralized markets, trading on margin, staking², etc. are all heavily influenced by, if not carbon copies of, traditional finance devices.

A common thread across many of these papers is the need for additional regulation to ensure market participants are provided the same protections as traditional markets. Unfortunately, before any such additional regulation can be added, the role of existing regulation within cryptocurrency markets must be defined. While the regulation regarding such manipulations in traditional markets is circumscribed, this is rarely the case for cryptocurrency markets. Anderson provides an example of such a challenge as it applies to front-running [12]. Anderson outlines flaws in the language of existing law, and the incompatibility with the structure of cryptocurrencies and blockchain. Weaver holds an opposing opinion that, at least newly issued cryptocurrencies, largely fall squarely within existing regulation [13]. They both agree, however, that additional regulation is paramount.

In the absence of effective regulation, successful manipulations run rampant. Eskandari et al. [14] provides a high-level summary of manipulations as it pertains to smart contracts and decentralized exchanges with the intention of bringing awareness to this type of manipulation. Daian et al. [15] take this research further, showing specifically how this process works on the Ethereum blockchain through the use of arbitrage bots, and the potential profit opportunities available to such manipulators. Not only do they find that it is possible, but that it is a present threat to blockchains. Piet et al. build on this by providing a method to identify actual front-running opportunities in historical blockchain data, as well as identifying extreme profit opportunities for arbitrage attacks [16]. However, arbitrage and manipulation are not isolated to decentralized exchanges or smart contracts. Czapliński et al. find that arbitrage can be profitable on exchanges by taking advantage of inefficient markets and trading between fiat currency pairs [17].

Existing literature has highlighted the prevalence of manipulations within the cryptocurrency ecosystem, as well as the arbitrage opportunities present at exchanges. Our work differs from this literature by investigating the arbitrage opportunities created by these market manipulations, as well as unexpected events.

2. Staking is, in essence, high-yield savings accounts without insurance.

3. Data and Methodology

We now describe our data sources, including the interruption events analyzed and the trade data collected, as well as our methodology for measuring the impact of these interruptions. We explain the different types of interruptions, how we use price, trade, and volume data to measure impact, and how we determine arbitrage opportunities.

3.1. Interruption Events

Events at cryptocurrency exchanges are commonplace, with some major ones making global headlines [18]–[20]. However, most events are shrouded in mystery, falling short of corroboration, or even consensus, among users. These events range from individual users claiming inability to access their funds, to hackers stealing millions of dollars worth of cryptocurrencies, to entire exchanges shutting down overnight. For this paper, we are interested in one particular subset of events: exchange interruption events. These events are temporary reductions in performance or accessibility, unique to individual exchanges.

To ensure the interruption events we investigate are legitimate, we only consider events which are corroborated by the targeted exchange itself. One of the most popular exchanges, Bitfinex, has a status page dedicated solely to self-reported events [21]. Each event includes a brief description of the event, the time it was posted to the status page, intermediate updates, and the time at which the issue was resolved. We collected all of the events, filtered out any events not related to reduced performance or accessibility, and grouped them into three types: DDoS, delay, and outage.

An event is considered a “DDoS” when Bitfinex explicitly cites this as the root cause of the reduced performance. “Delay” events are defined as any unexpected reduction in accessibility or performance without DDoS being explicitly cited as the primary cause. Finally, “outage” events describe an unexpected *complete* loss in access to the exchange or related services. Whether that outage does in fact lead to a complete loss of access is immaterial. As we will demonstrate, some outages continue to permit some trades during the reported event. These events, and their duration, can be seen in Figure 1.

In total we collected 40 interruptions: 12 DDoS, 14 delay, and 14 outage. The duration of the event is the time elapsed between the first official announcement from Bitfinex, and the final resolution update. The three longest duration events were all DDoS’s, with the longest one lasting seven days. Some events were extremely short-lived, with the shortest being an 8 minute outage in January 2019.

3.2. Event Types

Each type of interruption has unique implications. A DDoS attack, for example, requires the existence of a malicious actor initiating the attack, and therefore a motivation for the attack. Existing literature, such as the overview done by Dragomiretskiy [22] details a variety of motivations for DDoS attacks. One of the motivations provided by Dragomiretskiy




Figure 1: Timeline of exchange interruptions at Bitfinex, with their duration plotted on the y-axis in log scale.

is “manipulation”, that is, malicious actors use these attacks to produce abnormal market conditions from which they can obtain a profit. Abhishta et al. [23] investigate abnormal returns in volume at Bitfinex surrounding these events, and find that there are significant deviations, although it often recovers within a single day. While the impact suffered by the exchange is quantified in both papers, the motivation behind the attacks is not a primary focus. Prior to these reports, Vasek et al. [24] had performed a case study of Mt. Gox, noting the possible relationship between DDoS attacks, volume, and exchange rate.

The source of outage and delay events are not as explicit. While these could very well be DDoS attacks too, Bitfinex does not explicitly say so, thus malicious intent cannot be ascribed. Bitfinex relies on several third parties to operate the exchange, each susceptible to their own troubles. However, outages are unique as it relates to user access. While DDoS attacks and delays often impact only a portion of the users or services, outages are universal. No users can access the exchange, and no service is left online.

3.3. Bitcoin Trade Data

To measure the impact of these interruptions, we collect average price, total volume, and trade count data, which Bitfinex offers at five minute granularity. We then compare the Bitfinex data to aggregate data representative of the global bitcoin market. This data is collected from CoinMarketCap [25], also at five minute granularity. At the time of writing, CoinMarketCap tracks 569 exchanges, and provides aggregate price and volume data.

We collect all five minute data between January 1, 2015, and December 31, 2022. During August 2016, Bitfinex suffered a breach and shut down the exchange for several days.

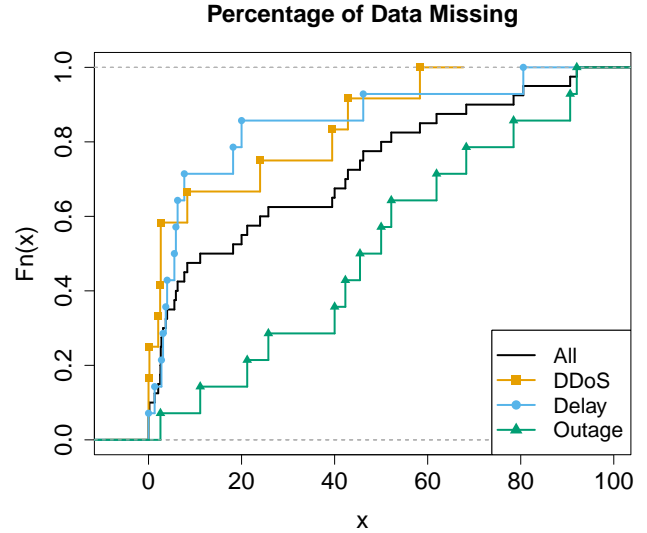


Figure 2: All events (black) as a function of the percent data missing during an event. Denial of Service events (orange), delays (green), and outages (blue) are also shown individually.

Therefore, we have no data from Bitfinex for approximately seven days between August 3, 2016 and August 10, 2016.

3.4. Arbitrage Opportunities

A bitcoin price difference between Bitfinex and CoinMarketCap does not automatically warrant an arbitrage opportunity. The arbitrage process itself incurs costs which must be outweighed by the profit before arbitrage is a viable strategy. These costs include transaction fees, withdrawal and deposit fees, etc. The arbitrage process for bitcoin would involve the purchase of bitcoin at an exchange with a discounted price. The arbitrageur would then need to move this bitcoin, using the blockchain, to the off-loading exchange. Once the bitcoin is at this second exchange, the bitcoin needs to be sold. The threshold defined by this assortment of costs is often referred to as the arbitrage band. We use the following equation to calculate the arbitrage band of bitcoin across centralized exchanges:

$$Arb = \pm((Tx + W + Tr * 2) * 2) \quad (1)$$

Here, Arb is the arbitrage band price, Tx is the cost of sending a transaction on the bitcoin network between exchanges, W is the fee for withdrawing bitcoin from the initial exchange, and Tr is the trading fee. The trading fee will be incurred twice, once when bitcoin is purchased at the initial exchange, and again when it is sold at the destination exchange. There is occasionally a deposit fee, however it is extremely rare, and often abandoned by exchanges shortly after adopting it. As a result, we do not factor it into our equation. The arbitrage band is two values, the positive and negative result of the equation, because arbitrage can occur both into, and out of, an exchange. The transaction fee, Tx , can be seen from any blockchain explorer, and we use daily




Figure 3: Percent difference between event trades and volume. Events are divided into quartiles (top row) based on percentage of missing data (second row) during the interruption. The data is further divided within each quartile by their relation to the interruption: Before (B), During (D) or After (A).

averages collected from YCharts [26]. The withdrawal fee, W , and the trading fee, Tr , vary between exchanges and over time. Upon investigating several of the most popular BTC exchanges, we observe an average withdrawal fee of 0.004 BTC, and a range of trading fees averaging 0.5%. To account for any other minor costs, as well as the risk involved in the arbitrage opportunity dissolving before the bitcoin can be off-loaded, we multiply this arbitrage band by two.

4. Analysis and Results

In this section we explain the analysis used to measure the impact of interruptions on the volume, trades, and price deviations of bitcoin on Bitfinex. We then quantify the arbitrage opportunities these interruptions facilitate.

4.1. Impact on Trading Activity

The events reported by Bitfinex, in addition to the broad type classification, are accompanied by a brief description. The severity and duration varies between events, even within the same type. Therefore, we first quantify the impact of events in terms of how much trading data is missing between first official reporting of interrupted services, and the notice of services restored. Missing data is identified by five minute intervals missing from the data provided by Bitfinex. The missing data per event can be seen in Figure 2. For all events, we observe behavior similar to a step function, with several distinct groupings of missing data around their quartiles.

We group interruption events based on their quartile of missing data, and include six hours of data before and after to identify event specific changes in behavior. We then measure the percent difference from mean volume and trades within

Group	Before		During		After	
	Mean	Median	Mean	Median	Mean	Median
Q1	-4.06%	-4.17%	24.9%	24.9%	-14.4%	-18.1
	-3.77%	0.98%	15.3%	18%	-6.71%	-5.36%
Q2	21.1%	29%	3.62%	1.54%	-27.1%	-38.6%
	11.6%	11.1%	14.7%	-11.4%	-17.4%	-11%
Q3	-18.7%	-29.6%	44.6%	42.3%	4.93%	5.28%
	-8.23%	-15.1%	-17.5%	-23.2%	21.5%	21.2%
Q4	-15.4%	-26%	35.7%	-5.09%	13.6%	16.3%
	4.85%	7.96%	-71.2%	-73.4%	29.1%	24.2%
DDoS	-4.24%	-12.4%	15.7%	19.3%	3.58%	3.96%
	-2.72%	0.98%	-7.87%	-4.51%	10.4%	10.3%
Delay	1.3%	-9.92%	23.1%	3.79%	-12.8%	-28.9%
	3.05%	3.55%	1.54%	-15.2%	-5.6%	-10.1%
Outage	-12.1%	-21.9%	47.9%	37.8%	-8.57%	-13.5%
	2.42%	-2.11%	-42.8%	-48.8%	16.2%	12.5%

TABLE 1: Percent difference from mean by grouping and relation. The first four groups split the data into quartiles. The last three groups are categorized by event type. Within each group, the top row is the volume percent difference, and the bottom row is the trade percent difference. The columns are split into their relation, before, during, or after the event. Within each relation is the mean and median value.

each quartile based on their relation to the actual interruption (i.e. before, during, or after). These results are shown in the first eight rows of Table 1, and visualized in Figure 3. The first quartile, which is less than 3% missing data, shows an increase in both trades and volume during events. In contrast, the second quartile, between 3% and 11% missing data, experiences increased trades and volume before events, and decreased trades and volume after events. The third quartile, between 18% and 45% missing data, has increased volume during events, despite a reduction in trades. Finally, the fourth quartile, where we see 46% to 92% of data missing, the during event trades drop dramatically, while volume remains largely unaffected.

The CDF plot in Figure 2 also breaks down missing data by type. As expected, we see that outages account for a significantly larger amount of missing data. We group our data by type, again adding six hours of data before and after the interruption, and measure the difference in mean. The results can be seen in the last six rows of Table 1, and are visualized in Figure 4. When grouped by interruption type, we can see all types have increased volume during the interruption, despite stagnant or decreasing trades. That is, the volume to trade ratio increases during interruptions.

Large deviations in volume are not uncommon in the cryptocurrency space. Large price swings, the passing of cryptocurrency related legislation, and both positive and negative news headlines are only some examples of common factors influencing cryptocurrency trading behavior. In Figure 5 we can see an example of the volume difference behavior for a delay event on June 14, 2016. During the twenty minutes preceding the interruption, we see a dramatic decrease in volume at Bitfinex, while the global aggregate volume remains relatively constant. Shortly after the event begins, Bitfinex




Figure 4: Percent difference between overall event trades and volume. Events are grouped by interruption type. The data is further divided by their relation to the interruption: Before (B), During (D) or After (A).




Figure 6: Timeline of the price difference of bitcoin between Bitfinex and CoinMarketCap during the event window. The arbitrage band is plotted (blue dashed), as well as the interruption beginning (orange), and end (green). Periods beyond the arbitrage threshold are filled in red.




Figure 5: Five minute percent change in volume at Bitfinex (purple) and CoinMarketCap (blue) for a delay event in June, 2016. The orange line represents the first official notice of an interruption, and the green line represents the interruption being officially resolved.

volume increases for roughly an hour before finally stabilizing a couple hours prior to official recovery. The initial decrease in volume at the beginning of an event is expected, as the exchange suffers reduced performance. However, the subsequent increase in volume during the interruption is counter intuitive, although it matches the preliminary findings from Figure 4.

4.2. Arbitrage Opportunities

The existence of such disproportionately large trades *during* interruptions could be explained by arbitrage opportunities. We identify a valid arbitrage opportunity as any five minute interval when the difference between the price of bitcoin at Bitfinex and CoinMarketCap will exceeds the arbitrage band defined by Equation 1. Figure 6 shows the same June 14, 2016 delay event, with the arbitrage band plotted on top of the price difference. During this event, we first observe the price difference exceed the lower arbitrage band shortly before the beginning of the delay. In this case, bitcoin is being sold at a discount on Bitfinex. During the delay the price difference rebounds, surpassing the upper arbitrage band, and begins being sold at a premium on Bitfinex. The cause of the delay is resolved shortly thereafter, and the price difference quickly returns within the arbitrage bands. All 40 events are plotted in a similar fashion, and can be seen in the Appendix. The events are shown in chronological order, with the date and time of the event, as well as the interruption type, shown in the plot title. Several of the other events show a similar trough to peak behavior, while others appear to have a single peak or trough.

To determine how common this behavior is outside of specific events, we bin our entire data set from January 2015 to December 2022 into six hour segments, and label each segment as having a significant arbitrage opportunity if any five minute interval in the window experiences a deviation beyond the arbitrage band. The results can be seen in the first two rows of Table 2. In this case we see significant differences between event periods and non-event periods, with the former being nearly 46% more likely to occur at the same time as a significant arbitrage opportunity.

In addition to the increased prevalence of interruptions

Event	Arbitrage	Count		Proportion		χ^2
		True	False	True	False	
True		102	47	68.5%	31.5%	27.49 (1.58e-7)
False		4,316	4,898	46.8%	53.2%	

TABLE 2: Contingency table with proportions and chi-square results (p-value). Columns represent the arbitrage band being exceeded, rows represent events occurring.

Type	Mean	Median	Count
None	82.1	20	9,739
DDoS	109	20	79
Outage	81.4	40	46
Delay	94.5	40	31

TABLE 3: The count of continuous price difference runs above or below the arbitrage band, as well as the mean and median duration (in minutes) of these runs.

during arbitrage opportunities, we observe a supporting relationship between events and arbitrage opportunities. We perform run length encoding on each five minute interval based on whether the price difference exceeds the arbitrage band, allowing us to measure how long each deviation lasts. Table 3, breaks down the average duration of these deviations by event type. In the absence of any interruption, significant arbitrage opportunities persist for just over 82 minutes on average, with a median duration of 20 minutes. During a DDoS interruption, the average arbitrage opportunity lasts 33% longer, persisting for 109 minutes, although the median duration is the same. Deviations coinciding with delay interruptions last on average 12 minutes longer, with twice the median duration. Finally, outages are roughly the same duration as non-event deviations on average, but twice the median duration.

4.3. Profit as a Motivation

While the information provided by Bitfinex regarding the interruptions make it hard to identify their origin, the environment in which exchanges currently operate is ripe for criminal exploitation. A malicious actor, after observing a broad price movement of bitcoin, can easily set up automatic future buy or sell orders on an exchange. They can then order a DDoS attack on the exchange, preventing future modifications to the order book as the global price of bitcoin moves further into their favor. The automatic trades initiated by the malicious actor are settled amidst the disarray caused by the interruption. Once the profitable opportunity has passed, the malicious actor can cease the attack and begin off-loading their ill-gotten gains.

The potential profit of such a scheme is substantial. By multiplying each reported five minute volume at Bitfinex by the five minute price difference between exchanges, we can estimate the profit available to an arbitrageur. In Figure 7 we see a timeline of interruptions, with their potential profit from arbitrage plotted on the y-axis in log scale. Four events, two outages and two DDoS attacks, provided an arbitrage opportunity exceeding \$10 million. Five more events were profitable beyond \$1 million. The median profit for all events is \$64,821, and a total profit across all events of \$105,807,735.




Figure 7: Timeline of events with the potential profit from arbitrage plotted on the y-axis in logarithmic scale.

	Mean	Median	Max
Profit	\$2,580,676	\$64,821	\$32,793,927
Ratio	32.5%	28.7%	71.5%

TABLE 4: Total profit and profit to cost ratio of arbitrage opportunities during events.

These events are not without risk, however, as the profit obtainable through arbitrage is often a small portion of the overall cost. Due to the profit opportunity being bounded by the arbitrage band, a transaction which settles at or below this band is not profitable. If an arbitrageur fails to accurately predict the profit opportunity, they risk a substantial loss. In Table 4 we can see summary statistics about the ratio of potential profit after accounting for costs within the arbitrage band. For an arbitrageur seeking to make the median profit of \$64,821, the median principal they are risking is \$161,036.

5. Conclusion

In this paper we investigate the impact of interruption events at Bitfinex, and the arbitrage opportunities they facilitate. We observe intuitive decreases in volume preceding events, followed by unexpected spikes in volume during the interruption. We observe large deviations between the local price at Bitfinex and the larger market during these interruptions. Large price differences for the same asset in different markets present a profitable trading strategy for arbitrageurs. We find that for all instances of the arbitrage band being exceeded, it is 46% more likely to witness a simultaneous exchange interruption event, with 68.5% of events occurring during such a circumstance. Additionally, we find that the duration of these price deviations beyond the arbitrage band are longer during these interruptions than during non-event periods. For DDoS events, the average duration is 27 minutes longer, although the median duration is 20 minutes for both. Delay events are 12


minutes longer on average, and both delays and outages are twice the median length at 40 minutes. These interruptions provide sizable profit opportunities, with a median profit of \$64,821, and a total profit across all events of \$105,807,735.

Acknowledgements

As part of the open-report model followed by the Workshop on Attackers & CyberCrime Operations (WACCO), all the reviews for this paper are publicly available at: github.com/wacco-workshop/WACCO/tree/main/WACCO-2023.

References

- [1] N. Weaver, "Risks of cryptocurrencies," *Communications of the ACM*, vol. 61, no. 6, pp. 20–24, May 2018. [Online]. Available: <https://dl.acm.org/doi/10.1145/3208095>
- [2] A. Feder, N. Gandal, J. T. Hamrick, and T. Moore, "The impact of DDoS and other security shocks on Bitcoin currency exchanges: evidence from Mt. Gox," *Journal of Cybersecurity*, vol. 3, no. 2, pp. 137–144, Jun. 2017. [Online]. Available: <https://academic.oup.com/cybersecurity/article/3/2/137/4831474>
- [3] P. McCorry, M. Möser, and S. T. Ali, "Why Preventing a Cryptocurrency Exchange Heist Isn't Good Enough," in *Security Protocols XXVI*, V. Matyáš, P. Švenda, F. Stajano, B. Christianson, and J. Anderson, Eds. Cham: Springer International Publishing, 2018, vol. 11286, pp. 225–233, series Title: Lecture Notes in Computer Science.
- [4] "Market Operator Obligations Targeted Review – NZX." [Online]. Available: <https://www.fma.govt.nz/library/reports-and-papers/moo-targeted-review-nzx/>
- [5] "EXCLUSIVE-Nasdaq pricing system focus of SEC outage review -source," *Reuters*, Aug. 2013. [Online]. Available: <https://www.reuters.com/article/nasdaq-halt-probid-ukl2n0gs19a20130827>
- [6] S. Krückeberg and P. Scholz, "Decentralized Efficiency? Arbitrage in Bitcoin Markets," *Financial Analysts Journal*, vol. 76, no. 3, pp. 135–152, Jul. 2020. [Online]. Available: <https://www.tandfonline.com/doi/full/10.1080/0015198X.2020.1733902>
- [7] I. Makarov and A. Schoar, "Trading and arbitrage in cryptocurrency markets," *Journal of Financial Economics*, vol. 135, no. 2, pp. 293–319, Feb. 2020. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0304405X19301746>
- [8] F. Eigelshoven, A. Ullrich, and D. Parry, "Cryptocurrency Market Manipulation: A Systematic Literature Review," 2021.
- [9] D. Twomey and A. Mann, *Fraud and Manipulation within Cryptocurrency Markets*. John Wiley & Sons, Jun. 2020, google-Books-ID: EwXaDwAAQBAJ.
- [10] A. Feder, N. Gandal, J. Hamrick, T. Moore, and M. Vasek, "The Rise and Fall of Cryptocurrencies."
- [11] J. M. Griffin and A. Shams, "Is Bitcoin Really Untethered?" *The Journal of Finance*, vol. 75, no. 4, pp. 1913–1964, 2020, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12903>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12903>
- [12] J. P. Anderson, "Insider Trading and Cryptoassets: The Waters Just Got Muddier Essay," *Iowa Law Review Online*, vol. 104, pp. 120–129, 2019. [Online]. Available: <https://heinonline.org/HOL/P?h=hein.journals/iowalrb10&i=123>
- [13] N. Weaver, "The Death of Cryptocurrency: The Case for Regulation," Dec. 2022.
- [14] S. Eskandari, S. Moosavi, and J. Clark, "SoK: Transparent Dishonesty: front-running attacks on Blockchain," Apr. 2019, arXiv:1902.05164 [cs]. [Online]. Available: <http://arxiv.org/abs/1902.05164>
- [15] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability," in *2020 IEEE Symposium on Security and Privacy (SP)*, May 2020, pp. 910–927, iSSN: 2375-1207.
- [16] J. Piet, J. Fairuze, and N. Weaver, "Extracting Godl [sic] from the Salt Mines: Ethereum Miners Extracting Value," 2022.
- [17] T. Czaplinski and E. Nazmutdinova, "Using FIAT currencies to arbitrage on cryptocurrency exchanges," *Journal of International Studies*, vol. 12, no. 1, pp. 184–192, Mar. 2019. [Online]. Available: https://www.jois.eu/?498,en_using-fiat-currencies-to-arbitrage-on-cryptocurrency-exchanges
- [18] C. Porterfield, "FTX Collapse: A 'Substantial Amount' Of Assets Are Missing And May Have Been Stolen," section: Business. [Online]. Available: <https://www.forbes.com/sites/carlieporterfield/2022/11/22/ftx-collapse-a-substantial-amount-of-assets-are-missing-and-may-have-been-stolen/>
- [19] E. Livni, "Binance Blockchain Hit by \$570 Million Hack, Exposing Crypto Vulnerabilities," *The New York Times*, Oct. 2022. [Online]. Available: <https://www.nytimes.com/2022/10/07/business/binance-hack.html>
- [20] R. McMillan, "The Inside Story of Mt. Gox, Bitcoin's \$460 Million Disaster," *Wired*, section: tags. [Online]. Available: <https://www.wired.com/2014/03/bitcoin-exchange/>
- [21] "Bitfinex Status." [Online]. Available: <https://bitfinex.statuspage.io/>
- [22] S. Dragomiretskiy, "The influence of DDoS attacks on cryptocurrency exchanges."
- [23] A. Abhishta, R. Joosten, S. Dragomiretskiy, and L. J. Nieuwenhuis, "Impact of Successful DDoS Attacks on a Major Crypto-Currency Exchange," in *2019 27th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP)*, Feb. 2019, pp. 379–384, iSSN: 2377-5750.
- [24] M. Vasek, M. Thornton, and T. Moore, "Empirical Analysis of Denial-of-Service Attacks in the Bitcoin Ecosystem," in *Financial Cryptography and Data Security*, R. Böhme, M. Brenner, T. Moore, and M. Smith, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, vol. 8438, pp. 57–71, series Title: Lecture Notes in Computer Science.
- [25] "Cryptocurrency Prices, Charts And Market Capitalizations." [Online]. Available: <https://coinmarketcap.com/>
- [26] "Bitcoin Average Transaction Fee." [Online]. Available: https://ycharts.com/indicators/bitcoin_average_transaction_fee






Figure 8: Price difference between Bitfinex and CoinMarketCap for all 40 interruption events, as well as arbitrage bands, event start, and event end.