Price Manipulation in the Bitcoin Ecosystem

Neil Gandal[a], JT Hamrick[b], Tyler Moore[b,*], Tali Oberman[a]

[a] *Berglas School of Economics, Tel Aviv University*
[b] *Tandy School of Computer Science, The University of Tulsa*

Abstract

To its proponents, the cryptocurrency Bitcoin offers the potential to disrupt payment systems and traditional currencies, with around $30 billion of the coins in circulation today. It has also been subject to security breaches and wild price fluctuations. We identify and analyze the impact of suspicious trading activity on the Mt. Gox Bitcoin currency exchange, in which approximately 600,000 bitcoins (BTC) valued at $188 million were fraudulently acquired. During both periods, the USD-BTC exchange rate rose by an average of four percent on days when suspicious trades took place, compared to a slight decline on days without suspicious activity. Based on rigorous analysis with extensive robustness checks, we conclude that the suspicious trading activity likely caused the unprecedented spike in the USD-BTC exchange rate in late 2013, when the rate jumped from around $150 to more than $1,000 in two months.

Keywords: Bitcoin, cryptocurrencies, fraud, exchange rate manipulation
JEL classification: E42, E31, E39.

1. Introduction

Bitcoin has experienced a meteoric rise in popularity since its introduction in 2009 (**?**). While digital currencies were proposed as early as the 1980s, Bitcoin was the first to catch on. The total value of all bitcoins in circulation today is around $28 billion (**?**), and it has inspired scores of competing cryptocurrencies that follow a similar design. Bitcoin and most other cryptocurrencies do not require a central authority to validate and settle transactions.

*Corresponding Author.
Email addresses: `gandal@post.tau.ac.il` (Neil Gandal), `jth563@utulsa.edu` (JT Hamrick), `tyler-moore@utulsa.edu` (Tyler Moore), `otalika@yahoo.com` (Tali Oberman)

Instead, these currencies use only cryptography (and an internal incentive system) to control transactions, manage the supply, and prevent fraud. Payments are validated by a decentralized network. Once confirmed, all transactions are stored digitally and recorded in a public "blockchain," which can be thought of as an accounting system.

While bitcoin shows great promise to disrupt existing payment systems through innovations in its technical design, the Bitcoin ecosystem[1] has been a frequent target of attacks by financially-motivated criminals. In this paper, we leverage a unique and very detailed data set to examine suspicious trading activity that occurred over a ten-month period in 2013 on Mt. Gox, the leading Bitcoin currency exchange at the time.[2] We first quantify the extent of the suspicious trading activity and show that it constitutes a large fraction of trading on the days the activity occurred. We then examine whether and how this trading activity impacted Mt. Gox and the broader Bitcoin ecosystem.

Our main results are as follows:

- Prices rose on approximately 80 percent of the days that the suspicious trading activity occurred. By contrast, prices rose on approximately 55 percent of the days in which no suspicious trading activity occurred.

- During days with suspicious trades, on average, the USD/BTC exchange rate increased by approximately four to five percent a day. During the same period when no suspicious trades occurred, on average the exchange rate was flat to slightly decreasing.

- Trading volume increased substantially on days with suspicious trading activity, over and above the suspicious activity.

- The effects of rising exchange rates and increased trading volume were found not only on the Mt. Gox exchange where the suspicious trades took place, but also on the other leading currency exchanges.

[1]The Bitcoin ecosystem includes the core network for propagating transactions, the blockchain, and many intermediaries such as currency exchanges, mining pools and payment processors that facilitate trade. We use "Bitcoin" with a capital "B" to refer to the ecosystem and "bitcoin " with a small "b" or BTC to refer to the coin.

[2]See Appendix A for the market share of the cryptocurrency exchanges.

- The suspicious trading activity of a single actor was the likely cause of the massive spike in the USD/BTC exchange rate in which the rate rose from around $150 to over $1,000 in just two months in late 2013. The fall was nearly as precipitous: the Mt. Gox exchange folded due to insolvency in early 2014 and it has taken more than three years for bitcoin to match this rise.

1.1. Why Should We Care?

As this paper will show, the first time Bitcoin reached an exchange rate of more than $1,000, the rise was likely driven by manipulation. It took more than three years for these exchange rates to be reached again, and we are left to wonder whether the current spike was driven by legitimate interest or by something more nefarious. But, why should we care about possible price manipulation in bitcoin during 2013? After all, the Bitcoin ecosystem is not nearly as important as the New York Stock Exchange. Nonetheless, recent trends indicate that bitcoin is becoming an important online currency and payment system.

Additionally, the total market capitalization cryptocurrency assets has grown stunningly since the end of the period covered by our analysis. In January 2014, the market capitalization of all cryptocurrencies was approximately $14 Billion. As of September 2017, total market capitalization is approximately $145 Billion. That is a ten-fold increase.

In the case of bitcoin, during the one year period ending in mid-May 2017, the market capitalization increased massively, from around 7 Billion USD to 28 Billion USD (**?**). That is an increase of approximately 300 percent in one year. The market cap of other cryptocurrencies surged by even more. In the one year period ending in mid-May 2017, the market value of cryptocurrencies excluding bitcoin surged by more than 1,900 percent (**?**). Hence, cryptocurrencies are becoming more important. So it is important to understand how the Bitcoin ecosystem works or does not.

Further, despite the huge increase in market capitalization, similar to the bitcoin market in 2013 (the period we examine), markets for these other cryptocurrencies are very thin. The number of cryptocurrencies has increased from approximately 80 during the period we examined to 843 today! Many of these markets are thin and subject to price manipulation.

As mainstream finance invests in cryptocurrency assets and as countries take steps toward legalizing bitcoin as a payment system (as Japan did in April 2017), it is important to understand how susceptible cryptocurrency markets are to manipulation. Our study provides a first examination.

In terms of the macro-economic lessons, cryptocurrency manipulations tie in to a concern in trading in unregulated financial exchanges. The potential for manipulation in the Over-the-Counter (OTC) markets is a significant concern for financial regulators. OTC trading is conducted directly between two parties, without going through a stock exchange. In a recent white paper, the SEC noted that "OTC stocks are also frequent targets of market manipulation by fraudsters."[3] The SEC report also documents that OTC trading has increased significantly over time.[4]

For all of these reasons, we believe that it is important to understand how the Bitcoin ecosystem works and how it could be abused. In this paper, we have taken an initial step in that direction.

1.2. Road Map

The paper proceeds as follows. Section 2 discusses background and related work. In section 3, we explain our methodology for identifying the STA and detail evidence for why we deem these transactions suspicious. Sections 4 and 5 examine the data in detail, present our findings and show that our results are robust. Section 6 documents the potential for fraudulent trading in the cryptocurrency market today, while Section 7 concludes with further discussion.

2. Background and Related Work

Cryptocurrencies and associated markets represent a nascent but growing force within the financial sector. Bitcoin, which became the first popular decentralized cryptocurrency in 2009, is the most researched because it is the most successful of the digital currencies. Within the finance literature, there is growing interest in discovering what drives a "value-less"

[3]Outcomes of Investing in OTC Stocks, by Joshua White, December 16, 2016, U.S. Securities and Exchange Commission Division of Economic and Risk Analysis (DERA).

[4]In 2008 around 16 percent of U.S. stock trades were of the OTC type. By 2014, OTC trades accounted for forty percent of all stock trades in the US. Like cryptocurrency trading, OTC trades are not transparent and not regulated, and there is concern that such trading is more harmful than high-frequency trading via regulated exchanges (**?**).

currency. **?** investigate the bitcoin exchange rate in an effort to expand our understanding of the motivation behind the rise and fall of cryptocurrency values. **?** build a theoretical model to examine the exchange rate of virtual currencies. Additionally, **?** constructs a model for determining the value of a bitcoin-like cryptocurrency by calculating its cost of production. **?** concluded that investor attractiveness has had a significant impact on Bitcoin's price.[5] While the potential for manipulation to influence valuations is sometimes acknowledged, none of these papers considered how unauthorized trades like the ones we study could affect valuations.

Unregulated cryptocurrency exchanges, such as Mt. Gox, are an essential part of the Bitcoin ecosystem. For most users, it is through currency exchanges that bitcoins are first acquired. As exhibited by the rise and fall of Mt. Gox, no cryptocurrency exchange is too big to fail. As reported by **?**, by early 2013, 45% of Bitcoin exchanges had closed, and many of the remaining markets were subject to frequent outages and security breaches. **?** performed an in-depth investigation of denial-of-service attacks against cryptocurrency exchanges and other Bitcoin services, documenting 58 such attacks. **?** conducted the first econometric study of the impact of denial-of-service attacks on trading activity at Bitcoin exchanges, leveraging Vasek et al.'s data on attacks. They show that trading volume becomes less skewed (fewer large trades) the day after denial-of-service attacks targeted the Mt. Gox exchange. In this paper, we use the same trade data to identify unauthorized trading and examine the effects of such trading on the Bitcoin ecosystem.

Due to their relatively lawless nature, cryptocurrencies are under constant threat of attack. Numerous researchers have conducted studies in order to document and combat threats such as Ponzi schemes (**?**), money laundering (**?**), mining botnets (**?**), and the theft of "brain" wallets (**?**). **?** attempt to identify suspicious trading activity by building a graph of Bitcoin transactions found in the public ledger. **?** examine the blockchain to determine whether bitcoin transactions are are truly anonymous. They successfully link transactions back to popular Bitcoin service providers, such as currency exchanges. None of these papers can associate individual transactions with specific users at currency exchanges. Our data

[5]**?** examine competition among cryptocurrencies. They find that the data are consistent with strong network effects and winner-take-all dynamics.

includes the user ID. Hence, we can associate trades with particular users.

For a more complete review of the literature, see **?** for coverage of technical issues and **?** for a discussion of Bitcoin's design, risks and open challenges.

2.1. Related Work on Price Manipulation

The academic literature on price manipulations of stocks includes **?**; they examined U.S. Securities and Exchange Commission litigation against market manipulators in OTC markets. They find small, illiquid stocks are subject to manipulation and that stock prices, volume, and volatility increase during the alleged manipulation period, but end quickly once the scheme is over. They note "while manipulative activities seem to have declined on the main exchanges, it is still a serious issue in the over-the-counter (OTC) market in the United States." Many of the more than 800 cryptocurrencies available today are illiquid and are characterized by very low volumes on most days and volume and price spikes. **?** studied OTC companies that hire promoters to engage in secret stock promotions to increase their stock price and trading volume. They find that the "promotions" coincide with trading by insiders. **?** show that OTC stocks have lower levels of liquidity than a matched sample of similar NASDAQ-listed stocks.

3. Identifying Suspicious Trading Activity on Mt. Gox

3.1. Exchange Activity

In early 2014, in the midst of theft allegations, the Mt. Gox transaction history was leaked. The Mt. Gox data dump gave access to approximately 18 million matching buy and sell transactions which span April 2011 to November 2013. These data are much more finely grained than data we would be able to get from the blockchain or public APIs for two reasons. First, a majority of the trading activity is recorded only by the exchange. Second, the exchange links transactions by the user account.

Data from the dump include fields such as transaction ID, amount, time, currency, and user country and state codes. Also included is the user ID, which is the internal number associated with Mt. Gox users. The user ID is crucial as it enables us to link transactions by the same actor.

We supplemented the Mt. Gox data with publicly available daily aggregate values from `bitcoincharts.com`. This data was used to verify trading volumes, to compare Mt. Gox exchange rates to other leading platforms, and to verify daily USD to BTC exchange rates. We discuss how we built the dataset in detail in Appendix B.

3.2. Suspicious Trading Activity

In early 2014, after the Mt. Gox data leak, several individuals trading on Mt. Gox found what they considered "suspicious activity" and wrote extensively about it (**??**). We conducted our own analysis of the data, confirming much of what was reported on the blogs.[6] In Appendix B, we discuss why this trading activity should be deemed suspicious, along with a description of the behavior. We carefully go through the details that confirmed that the relevant transactions were suspicious. Here we present a brief description of the trading activity and what effect it had on the markets. We use the names given by the blogs to the suspicious traders: (1) the"Markus bot" and (2) the "Willy bot."

3.2.1. Suspicious Trader 1: the Markus Bot

Markus began "buying" bitcoin on 2013-02-14 and was active until 2013-09-27. His account was fraudulently credited with claimed bitcoins that almost certainly were not backed by real coins. Furthermore, because transactions were duplicated, no legitimate Mt. Gox customer received the currency Markus supposedly paid to acquire these claimed coins. On 33 of the 225 days the account was active, Markus acquired 335,898 bitcoins (worth around $76 million).

3.2.2. Suspicious Trader 2: The Willy Bot

Unlike Markus, Willy did not use a single ID; instead, it was a collection of 49 separate accounts that each rapidly bought exactly 2.5 million USD in sequential order and never sold the acquired bitcoin. The first Willy account became active on 2013-09-27, a mere 7 hours and 25 minutes after Markus became permanently inactive, and we are able to track Willy activity until our data cutoff on 2013-11-30. Each account proceeded to spend exactly 2.5 million USD before becoming inactive. Then the next account would become active and the

[6] Online commentary about these trades frequently refer to the traders as 'bots' (e.g., (**??**)).

process would repeat. Unlike Markus, it appears that Willy was interacting with real users. While accounts of these users were "nominally" credited with fiat currency, Willy likely did not pay for the bitcoins.

Willy traded on 50 of the 65 days before the data cutoff. In total, Willy acquired 268,132 bitcoin, nominally for around $112 million. While Willy acquired slightly fewer bitcoins than Markus, the Markus activity occurred on 33 days spread over a 225-day period. Thus, the Willy activity was much more intense. Together, the bots acquired around 600,000 bitcoins by November 2013.

Recently, in a trial in Japan, the Former Mt. Gox, CEO Mark Karpeles, confirmed that the exchange itself operated the "Willy" accounts and that the trades were issued automatically (?).[7]

What motivated the operation of these bots?. The publicly reported trading volume at Mt. Gox included the fraudulent transactions, thereby signaling to the market that heavy trading activity was taking place. Indeed, we later show that even if we set aside the fraudulent activity, average trading volume on all major exchanges trading bitcoins and USD was much higher on days the bots were active. The associated increase in "non-bot" trading was, of course, profitable for Mt. Gox, since it collected transaction fees.

But the Willy Bot likely served another purpose as well. A theory, initially espoused in a Reddit post shortly after Mt. Gox's collapse (?), is that hackers stole a huge number (approximately 650,000) of bitcoins from Mt. Gox in June 2011 and that the exchange owner Mark Karpales took extraordinary steps to cover up the loss for several years.[8]

Note that Bitcoin currency exchanges function in many ways like banks. Customers buy and sell bitcoins, but typically maintain balances of both fiat currencies and bitcoins on the exchange without retaining direct access to the currency. If Mt. Gox was trying to hide the absence of a huge number of BTC from its coffers, it could succeed so long as customers remained confident in the exchange. By offering to buy large numbers of bitcoins, Willy

[7] It also appears that Karpeles operated the Markus Bot as well, and this is where most of the prosecutor's evidence against Karpales has focused.

[8] When Mt. Gox folded, it initially announced that around 850,000 bitcoins belonging to customers and the company were missing and likely stolen. Shortly thereafter, Mt. Gox found an additional 200,000 bitcoins. Hence, the total loss was 650,000 bitcoins.

could prop up the trading volume at Mt. Gox and "convert" consumer "bitcoin" balances to fiat money. This could work, i.e., stave off collapse of the exchange, as long as users who sold bitcoin had enough confidence to leave the bulk of their fiat balance at Mt. Gox. If consumers wanted to take out bitcoins, Mt. Gox would immediately have to supply them. On the other hand, if consumers wanted to redeem the fiat cash in their accounts, Mt. Gox could delay the withdrawal by saying that their bank was placing limits on how much fiat cash Mt. Gox could withdraw in a particular period. This indeed happened, and some (but not all) consumers could not withdraw cash from their fiat accounts during the last couple of months before Mt. Gox shut down. By using this strategy, the Willy bot could turn the Mt. Gox' "bitcoin deficit" into a fiat currency deficit. This may have delayed the inevitable crash of Mt. Gox. Although this cannot work in the long-term, Bernie Madoff, a once respected stockbroker, kept a similar scheme running for many years.

4. Impact of Suspicious Purchases: Preliminary Analysis

[Figure 1 about here.]

Figure 1 shows that the USD/BTC exchange rate increased dramatically during the period Willy was active. We are, of course, not the first to notice that. But that in itself does not mean that Willy's activity *caused* the price rise. In this section and the next, we provide compelling evidence that the fraudulent activity likely *caused* the price rise. We first examine the impact on trading volume and then prices.

[Table 1 about here.]

4.1. Suspicious Purchases and Trade Volume

On the days they were active, Markus and Willy purchased large amounts of bitcoins. As Table 1 shows, Markus purchased on average $9,302$ BTC, which accounted for approximately 21 percent of Mt.Gox's daily volume of trades. On the days Willy was active, he purchased on average $4,962$ BTC, which accounted for 18 percent of Mt. Gox's daily volume of trades. Figure 2 gives a more detailed breakdown. It shows the fraction of daily BTC traded on the Mt. Gox exchange platform that were carried out by Markus and Willy, respectively.

The share of total trading volume remains significant, even taking into account trades on other platforms. Markus accounted for 12 percent and Willy 6 percent of the total trade on the four main exchanges trading bitcoin and USD on the days they were active. In addition to Mt. Gox, the other main exchanges trading USD/BTC during this time period were Bitstamp, Bitfinex and BTC-e. These exchanges accounted for more than 80 percent of the trading activity in BTC/USD during the period we analyze.

[Figure 2 about here.]

[Table 2 about here.]

We divided the data into four equal three-month periods, starting from December 1, 2012 (2.5 months before Markus was active) and ending when the leaked Mt.Gox dataset ends at the end of November 2013. The bulk of Markus's trades occur in period 3, while all of Willy's take place in period 4.

The increase in total trading volume cannot be accounted for by the rogue trades alone. Both Markus' Willy's activity were associated with much higher trading volume above and beyond their own contributions. On the days these actors were purchasing bitcoins, total volume on Mt. Gox and the other leading exchanges was significantly higher than on days when these bots were not active. Table 2 shows that during the 50 days Willy was active in period 4, he "purchased" approximately 3,900 bitcoins per day on Mt. Gox. Total median daily volume on Mt. Gox during these 50 days was approximately 26,000 bitcoins. During the 41 days that Willy was not active in the period, median daily volume on Mt. Gox was approximately 10,500 bitcoins. The differences in volume are similar across the other three platforms as well. Median total volume on the four exchanges was approximately 83,000 bitcoins on days Willy was active versus approximately 29,500 on days Willy was not active.

The same holds true for days that Markus was active in period 3. On the days that Markus was active during period 3 he "purchased" approximately 8,900 bitcoins per day on Mt.Gox. The total median daily volume on Mt.Gox on the days he was active in this period was 42,000 bitcoins, but only 17,400 bitcoins on the days he was not. The differences in volume are similar across the other three platforms as well. Median total volume on the four exchanges was approximately 68,000 bitcoins on days Markus was active in period 3 versus approximately 31,000 on days Markus was not active in period three. (See Table 2.) For a full breakdown of volumes on individual exchanges, see the tables in Appendix C.

Hence, although these bots differed in their method of operation, days in which either was active were associated with very high volume beyond the bots' direct contributions. It is likely their activity sent a signal to the market and encouraged others to enter and purchase bitcoins. This may be one of the reasons why their activity could have such a large effect on the bitcoin price. We conduct a preliminary examination of their effect on prices in the next section.

4.2. Suspicious Purchases and Price Changes: Preliminary Analysis

We would expect an association between the suspicious purchases and a rise in prices on Mt. Gox (and other exchanges as well.) This is because an upward shift in demand should lead to a rise in price. Although the activity took place exclusively on Mt. Gox, we are also interested in examining how it affected the other exchanges in the Bitcoin ecosystem.

On the days that there was suspicious trading activity on Mt. Gox, the descriptive evidence suggests that prices also tended to rise. The lines in the Figure 2 are colored green if the exchange rate rose and red if the exchange rate fell. We then examined whether the price changes differed on the days in which the fraudulent activity occurred. We did this for the 9.5 months Markus and Willy were active (and for which we have data) and observed how often the exchange rate rose on Mt. Gox, as indicated in Table 3. We can see that on days without suspicious activity, 55% of the time the exchange rate did in fact rise. But on the 82 days that there was suspicious purchasing activity, 79% of the time the exchange rate rose. According to a chi-squared test of proportions, it is unlikely that this difference was due to randomness ($p < 0.05$). This is preliminary evidence that this activity was associated with the price rise on Mt. Gox.

[Table 3 about here.]

Not surprisingly, similar patterns of price appreciation took place at other exchanges during this period. As shown in Appendix C, on days without unauthorized activity, the exchange rate on Bitstamp rose 55% of the time. However, on the 82 days that Markus or Willy acquired bitcoins, the exchange rate rose more than 80 percent of the time. This suggests that the suspicious trading on Mt. Gox spilled over to other exchanges. This makes sense because all of these platforms traded the same USD-BTC currency pair.

[Table 4 about here.]

Table 4 shows the percent of days in each period, in which there was suspicious trading activity. Markus was active over 8 months, which span over 4 periods. However, he was primarily active in period 3. Willy on the other hand was active for less than three months and all of the activity occurred in period 4. We have no data on any unauthorized activity from the end of period 4. Mt. Gox shut down shortly thereafter.

[Table 5 about here.]

In Table 5 we show how the daily movement in the exchange rate (closing price less opening price) changed, on average, on four main exchange platforms.[9] Since fraudulent activity essentially only occurred in the third and fourth periods, we focus on these two periods. Periods one and two can be viewed as benchmarks.

In period 3, when Markus' activity peaked, we don't see much change overall in the daily exchange rate. However, if we look at the days Markus is active, the average daily price increase is higher. This is true, both on Mt. Gox and on all the other platforms too.

In period 4, the sole period in which Willy was active, we see a big jump in the average daily exchange rate change. Separating the days on which Willy was active from those he was not, reveals a dramatic difference: In the case of Mt. Gox, the average USD/BTC rate increased by $21.85 on the 50 days Willy was active; it actually fell (by $0.88 on average) on days when Willy was not active. The same dramatic difference holds for the other exchanges as well.

We then look at daily 'return, which is the typical measure for assessing the performance of assets. Daily returns are defined to be the percentage change in the daily exchange rate, i.e., the closing price less the opening divided by the opening price. Table 5 also shows

[9] There is 24 hour trading, so the closing rate on one day is exactly the same as the opening rate on the following day. Bitfinex has fewer observations as it was not active until April, 2013.

the daily returns (in parentheses) for the four periods for days that Willy and Markus were active and days that they were not active. The table shows that the average daily returns when Markus was active in period 3 (which was his peak activity period) ranged from 1.9-2.9 percent on all four exchanges. On other days, the average return was slightly negative or all four exchanges.

Similarly, table 5 shows the daily returns (in parentheses) that the average daily returns when Willy was active (period 4) ranged from 4.8-5.0 percent on all four exchanges. On other days, the average return was slightly negative on all four exchanges.

These results are striking and make it very clear that the suspicious purchasing activity could have caused the huge price increases. The average daily returns when Markus was active were somewhat smaller than when Willy was active, but these daily rates of return appear non-trivial as well. In the following section, we will run regressions to control for other possible effects on the exchange rate.

5. Regression Analysis

The analysis in the previous section provides strong evidence that the suspicious activity on Mt. Gox may have affected prices on all exchanges. In this section, we use regression analysis to control for other events (like distributed denial of service (DDoS) attacks) that may have caused the changes in the exchange rate. We run regressions with the dependent variables being (I) the absolute daily price changes and (II) the daily returns on all four exchanges.

5.1. Daily Price Changes

We run the following regressions:

$$RateChange_t = \beta_0 + \beta_1 Markus_t + \beta_2 Willy_t + \beta_3 DDoS_t + \beta_4 DayAfterDDoS_t + \beta_5 Other_t + \epsilon_t \tag{1}$$

$$Returns_t = \beta_0 + \beta_1 Markus_t + \beta_2 Willy_t + \beta_3 DDoS_t + \beta_4 DayAfterDDoS_t + \beta_5 Other_t + \epsilon_t \tag{2}$$

Our first dependent variable, *RateChange*, is the daily difference in the exchange rate of BTC, i.e. the daily difference between the closing and opening prices.[10] Our second dependent variable, *Returns*, is the daily difference in the exchange rate of BTC, i.e. the daily difference between the closing and opening prices

We now describe our independent variables. *Markus* is a dummy variable that takes on the value one on the days Markus is active as a buyer. Similarly, we define the dummy variable *Willy*. *DDoS* is a dummy variable that takes on the value one on days a DDoS attack on Mt. Gox occurred. *Day after DDoS* is a dummy variable that takes on the value one on the day after a DDoS attack on Mt. Gox occurred. The variable *Other* (or *OtherAttacks*) is a dummy variable that takes on the value one on days that non DDoS attacks occurred.[11] ϵ_t is a white noise error term.[12] The subscript "t" refers to time. We have a total of 365 observations, except for Bitfinex which was not operating during period period one.

Equations (1) and (2) are reduced-form regressions. That is, we are not estimating demand or supply, but rather the effect of changes in exogenous right-hand-side variables on the endogenous variables (the daily rate change and the daily returns in percentage terms.) But in our case, the coefficients from these reduced form regressions are exactly what we want to measure. Summary statistics (and all other tables not in the text) appear in Appendix C.

[Table 6 about here.]

The results in Table 6 show that the coefficient representing Willy's activity is positive and significant: hence there is a very strong positive association between activity by Willy and prices on Mt. Gox. This regression confirms the striking results of Section 4. The estimated coefficient on the "dummy" variable for Willy is \$21.65, while the "estimate" in section 4 was \$21.85. This again suggests that the USD/BTC exchange rate rose on Mt. Gox

[10] Recall that closing prices on day t equal opening prices of day $t+1$ since there is 24 hour trading. The opening/closing price is at 24:00 GMT.

[11] Perhaps because it was the leading exchange during the period of our data, most of the DDoS attacks were on Mt. Gox.

[12] We check for autocorrelation of errors by calculating the Durbin Watson (DW) statistic for each regression. The value of DW is not statistically different from two in any of the four cases; this strongly suggests that there is no autocorrelation and a white noise error term is appropriate.

by more than 20 dollars a day on average on the days that Willy was active. The regressions for the other exchanges in the same table shows that price on that exchange also rose by 19-20 dollars a day on average on the days that Willy was active. Again the estimated coefficients are consistent with the "estimates" from the summary statistics in section 4.[13]

The estimated coefficient on the dummy variable representing Willy's activity is the only coefficient which is significant. Notably, denial-of-service attacks and other shocks do not appear to influence the exchange rate. While this does not conclusively prove that Willy's activity caused the price rise, it suggests that it was the likely cause of the significant price rise in the price of Bitcoin during this period. The estimated coefficient associated with Markus's activity is positive, but not significant, suggesting that Markus' more diffused activity was not associated with a large rise in the daily change (in levels) of the USD-BTC exchange rate.

[Table 7 about here.]

5.2. Daily Percentage Returns

Typically, in the finance literature, researchers examine daily returns to currencies in percentage terms, that is closing price less opening price divided by opening price. Hence, we now repeat the same exercise using daily percentage returns as the dependent variable, and employ the same independent variables as in the previous regressions.[14]

Table 7 shows that activities of the two bots led to similar rates of returns and that these returns were significantly higher than the returns earned during days in which the bots were not active. On days in which the bots were not active, the average rate of return was less than one percent (as the estimates of the constant show.) From the coefficients in Table 7, in the case of Willy, the daily returns across all exchanges were in a fairly tight range, ranging from 4.1 to 4.7 percent more when Willy was active than when he was not active. (On days when the suspicious actors were not active, there was no percentage change in the exchange rate.) All of the "Willy"coefficient estimates are significant at the 99% level of confidence.

[13]Controlling for other factors, the price change on days when the bots were not active was essentially zero, as the estimates of the constant show.

[14]We obtain virtually identical result using the natural log of returns i.e., the natural log of the closing price divided by the opening price.

In the case of Markus, the estimated coefficients in Table 7 show that the daily returns across the exchanges ranged from 2.7-4.3 percent more than when Markus was not active. The rates are similar to those when Willy was active. With the exception of Bitfinex, the "Markus" coefficient estimates are significant at the 99% level of confidence.[15]

6. Testing for Potential Price Manipulation Today

? describe one of the cases that involved price manipulation of "penny stocks." In that case, according to the SEC, the defendant placed purchase orders in small blocks at successively rising prices. The SEC alleged that this was part of a manipulative scheme to create the artificial appearance of demand for the securities in question, enabling unidentified sellers to profit and inducing others to buy these stocks based on unexplained increases in the volume and price of the shares."

Intentionally or not, this method resembles the one employed by the Markus and Willy bots. This suggests that one way to examine whether such price manipulation exists is to follow individual trades over time for each cryptocurrency - and see whether a pattern of systematic buying over time has occurred and whether such buying was associated with an increase in price. In order to control for periods of high demand for cryptocurrencies in general, one can compare these buying patterns with trends in bitcoin, the leading cryptocurrency.

[Table 8 about here.]

Researchers can use publicly available data on trading volume and price to raise red flags regarding possible price manipulation. As proof of concept (to examine the effects of increased trading volume on the price of cryptocurrencies,) we gathered publicly available data from coinmarketcap.com. These data give us access to cryptocurrencies tracked by the platform, which is an extensive though incomplete list. The data include daily market cap, trading volume and the open, high, low, and close price in USD for all currencies tracked. Starting from a total of 843 publicly traded currencies and 477,039 daily summaries for

[15] In the case of Bitfinex, the estimated coefficient on Markus' activity is 2.7, which is significant at the 10 percent level of confidence. Recall that the Bitfinex exchange was not operating in period one.

those cryptocurrencies, we sought to identify circumstances that might resemble the effects of fraudulent trades found in this paper.

We look for two patterns: first, a substantial market capitalization where profits could be made but thin enough for fraud to succeed; and second, a spike in daily trading volume that might drive returns higher. On the first count, we identified the 308 currencies which had a maximum market capitalization between \$1-100 million. On the second count, we compared the daily volume of each cryptocurrency to the average daily volume for that month and computed summary statistics for two overlapping groups. The first group consists of coins whose daily trading volume increased by at least 150% of the average daily trading volume for that month (e.g., the coin's trading volume jumped to \$2.5 million from a daily average of \$1 million). The second group considers more extreme jumps of at least 200% compared to that month's average trading volume. The reason we seek out these volume spikes is that we observed in Section 4.1 that the trading volume jumped over 200% on days when the bots were active.

As shown in Table 8, the first group (150%) consists of 19,212 events for 304 unique currencies. On the days when trading volume spiked, the coin's USD exchange rate increased by 26.8% on average (1.5% median.) By contrast, when the volume did not jump, the average price increase was 8.6% (median 0%).

For the second group requiring a 200% jump, the difference is even more stark. On days with volume spikes, the average price increase was 30.5% (median 2%), compared to an average price increase of 8.8% (0% median) on other days. While these jumps in trading volume and prices could certainly have an innocuous explanation, they nonetheless demonstrate the potential for fraud in a very opaque and unregulated market.

7. Concluding Remarks

In this paper, we used trade data delineated by user to conclude that the suspicious trading activity on the Mt. Gox exchange was highly correlated with the rise in the price of Bitcoin during the period we study. If the bot activity was indeed the cause, we have shown that manipulations can have important real effects. The suspicious trading activity of two actors were associated with a daily 4% rise in the price, which in the case of the second actor combined to result in a massive spike in the USD-BTC exchange rate from around \$150 to over \$1 000 in late 2013. The fall was even more dramatic and rapid, and it has taken more than three years for Bitcoin to match the rise during this period.

Given the recent meteoric rise in bitcoin to levels beyond the peak 2013 (and the huge increase in the prices of other cryptocurrencies), it is important for the exchanges to ensure that there is not fraudulent trading. The potential for manipulation has grown despite the increase in total market capitalization because there has been a very large increase in the number of cryptocurrencies. Currently, there are more than 300 cryptocurrencies with market capitalization between \$1 Million and \$100 Million. In January 2014, there were less than 30 coins with market capitalization between \$1 million and \$100 million. Hence, there are many more markets with relatively small market capitalization than there were in 2014. Thus, despite the 10-fold increase in market capitalization, the addition of so many "thin" markets in cryptocurrencies means that price manipulation remains quite feasible today. As shown in the prior section, these thin markets do exhibit sudden spikes in trading volume that drive the exchange rate upwards.

Since the Bitcoin ecosystem is currently unregulated, "self-policing" by the key players and organizations is essential. Further, as the Bitcoin ecosystem becomes more integrated into international finance and payment systems, regulators may want to reassess the policies that leave the ecosystem unregulated and take an active oversight role.

Acknowledgements

We are extremely grateful to the editor, Urban Jermann, and an anonymous referee; their comments and suggestions significantly improved the paper. We gratefully acknowledge support from research grants from the Interdisciplinary Cyber Research Center at Tel Aviv University and US National Science Foundation Award No. 1714291. We thank Maarten van Oordt and Nittai Bergman for very helpful and suggestions suggestions that greatly improved the paper. We also thank participants at the Central Bank Research Association conference at the Bank of Canada and the Workshop on the Economics of Information Security (WEIS) 2017 for helpful suggestions and comments.

Appendix A: Bitcoin Trading Market Share by Exchange

[Figure 3 about here.]

Appendix B: Dataset Validation and Details of Markus and Willy Activity

(i) Dataset Validation:

With the exception of a few key steps, validating the Mt. Gox data closely followed previous work done by **?** in which duplicates were removed by inspecting combinations of key fields. The duplicate rows contained matching values for user ID, time stamp, transaction type (buy/sell), and transaction amount. We examined two methods to remove duplicate entries. Both methods treated tuples as unique (user ID, timestamp, transaction type, amount in BTC, amount in JPY, i.e., Japanese Yen) with the more aggressive of the two methods excluding amount in JPY from the tuple.[16] Both methods produced results that were more consistent with other publicly available trading data than the original dataset. **?** chose to proceed with the less aggressive of the two strategies, which resulted in a clean dataset of approximately 14 million records. We chose the more aggressive method, but our results are robust to both methods.

During the data exploration phase, we discovered additional duplicate records that did not fit the unique tuple model outlined above. In these instances they appeared to be copies of either one side (buy/sell) of the transaction or of the entire transaction with minor alterations to the data in the "User_ID," "Money," and "Money_JPY" columns. The common factor used to start the removal process was the new user ID. We could find one side of the transaction by matching on this user ID, and then use the Money and Money_JPY columns to find the matching opposite side of the transaction. In total 5,991 additional rows were removed using this method, all involving a single user ID. We later identified these duplicate entries as originating from the trader denoted "Markus." We performed additional sanity checks of the data utilizing publicly available historical Mt. Gox trading data from `bitcoincharts.org`. We are confident that the data are high-quality.

(ii) Details of Markus and Willy Activity:

During initial data exploration we found a group of users with attributes that differed from the rest of the users in the dataset. In particular, for these users every transaction had "??" as an entry for the user country and user state fields. This appeared suspicious

[16]Mt Gox was based in Tokyo.

as these fields normally contain FIPS location codes, a NULL value, or "!!". One account containing the abnormal location values stood out when compared to the others because this account bought and sold bitcoins, where as the others only bought. We adhere to the naming convention in the blogs and refer to the first account as Markus.[17]

Upon closer inspection, Markus's trades raised many red flags. He never paid transaction fees and reportedly paid seemingly random prices for bitcoins. Most curious of all, we identified many duplicate transactions in which the amount paid was changed from an implausibly random price to one that was consistent with other trades that day.

Markus seemingly paid random rates for the bitcoins he acquired. For example, in two transactions that took place the same hour on 2013-04-03, he paid 0.000374 USD per bitcoin on one transaction and 925 489.67 USD per bitcoin on another.

Table 9 shows the wide range of rates that Markus paid. The table reports the number of purchases that Markus made for different ranges of rates. During the time when Markus traded, published exchange rates ranged from $20 to $229. Hence, any transactions with rates outside this range raise suspicion. In fact, only a quarter of Markus's trades fell within this range. 13% of the time, Markus paid less than one dollar, while in 821 transactions (3% of the total), he supposedly paid a rate of exceeding $100,000 per bitcoin!

[Table 9 about here.]

Upon closer inspection, the random exchange rates appear to come from transactions posted before Markus' transactions. Table 10 illustrates the pattern. Transaction 1362466144485228 was posted with user 238168 buying ≈0.398 bitcoin for 15.13 USD. Every Markus transaction that followed (indicated in bold) "borrowed" the Money, and Money_JPY values from the previous transaction. We confirmed this pattern of behavior throughout – whenever Markus bought, the amount paid came from a previous unrelated transaction, while the number of bitcoins acquired appears randomly.

[Table 10 about here.]

[17] Despite the fact that Markus sold bitcoin on a few occasions, most of his activity involved acquiring bitcoins.

Occasionally Markus would also sell bitcoin, and the BTC-fiat currency exchange rate for these transactions appears to be correct. For example, on 2013-06-02 Markus sold 31.5 bitcoins for 3 757.95 USD, or 119.3 USD per bitcoin, which is similar to the average rate paid by users that day. In total, Markus sold 35867.18 bitcoin worth approximately 4 018 681.65 USD in 2927 transactions on 6 different days.

As stated in section 3.2, we paid closer attention to what records to remove while de-duplicating the data. This was due to the fact that several transactions contained duplicate buy and sell rows; see Table 11 for an example of these transactions. We can see that apparently user 201601 sold one bitcoin twice at the same exact time, first to user 698630 for 15.13 USD and second to user 634 for 38.11 USD.

[Table 11 about here.]

Upon closer inspection, we concluded that the rows containing 15.13163 in the money columns are the original rows for this transaction. In every instance where duplicates were discovered they involved user 698630 and user 634; 634 appeared to "correct" the 698630. There are multiple oddities involving user 634. First, the numeric user ID is extremely low, which strongly suggests that it could be an internal Mt. Gox account. Second, prior to issuing the corrected transactions, user 634 bought and sold a total of 824,297.7 bitcoin between 2011-04-07 and 2012-08-01. This account was inactive for 197 days before we see it used again in the duplicate transactions involving Markus.

Table 12 summarizes the discrepancies between Markus's identities. 2 966 buy transactions made by 698630 were later duplicated as originating from user 634 at market prices. In total, as user 698630, Markus reportedly paid 1 080 617 USD for 67 452 bitcoin. When acting as user 634 instead, Markus "paid" 2 000 729 USD for the same transactions. This only includes the corrected transactions involving user 634; we ignore any trading activity that occurred before Markus was active. It is worth noting that only the amounts paid for bitcoins were altered, never the bitcoin amount. Additionally, for the 196 transactions where user 698630 sold bitcoin and we found a duplicate row with user 634, no monetary amounts were altered. Only the user ID had changed.

Finally, it is worth noting that the majority of transactions by user 698630 were never changed, despite the presence of often wild exchange rates. User 698630 only operated between February and September 2013, and during that time he purchased 268 446.09 BTC, reportedly at prices totaling $76.4 million. We note that this total USD amount should be viewed with caution, given that it is based on seemingly random exchange rates.

[Table 12 about here.]

In the case of Willy, in addition to the circumstantial evidence of sequential use and proximity to Markus, the most solid evidence we have that foul play was involved can be traced to the internal user ID. Previous research into the account IDs used for this activity showed that they were abnormally high for the time period in which they operated (**?**). Normal accounts for this time period had IDs that capped around 650000 where the users at the center of this research had IDs in the range of 658152-832432. Furthermore, several reports can be found online of the Mt. Gox trading API going offline for various periods of time in which no trading activity was being processed with one exception; Willy trading activity continued unabashed (**?**). On 2014-01-07 the trading API was offline for 90 minutes. During this time period the only activity being processed followed the exact buying pattern of Willy when he was active: 10-19 bitcoins purchased every 6-20 minutes.

Appendix C: Descriptive Statistics and Other Tables

[Table 13 about here.]

[Table 14 about here.]

[Table 15 about here.]

[Table 16 about here.]

[Table 17 about here.]

[Table 18 about here.]

[Table 19 about here.]

List of Figures

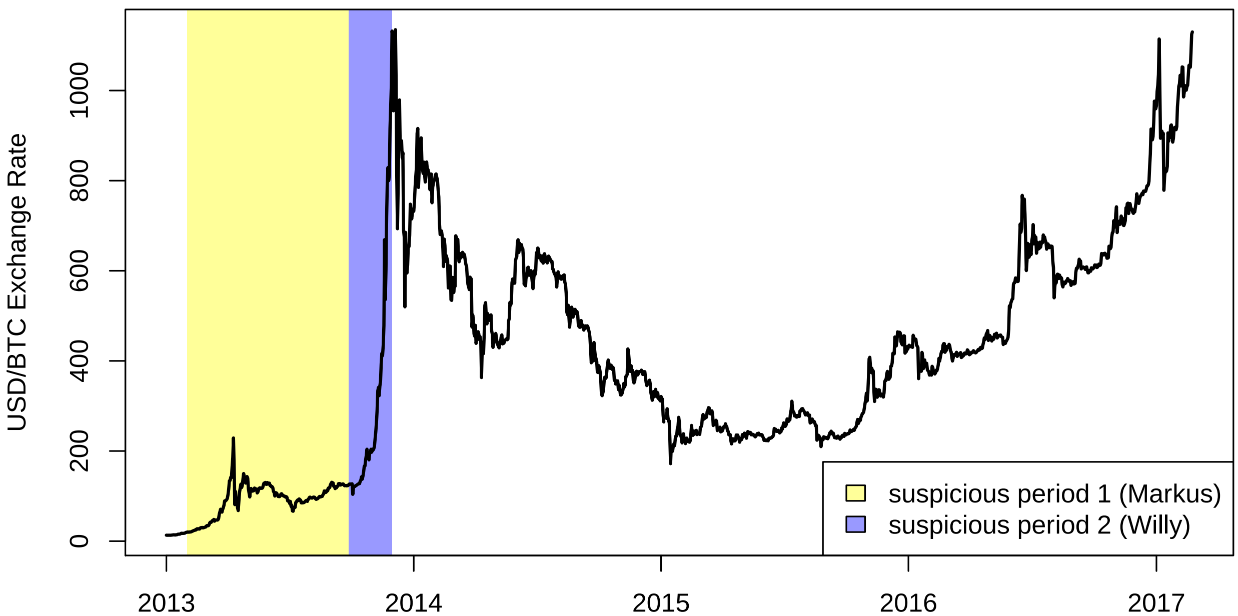

Figure 1: Bitcoin-USD exchange rate with periods of suspicious activity shaded.

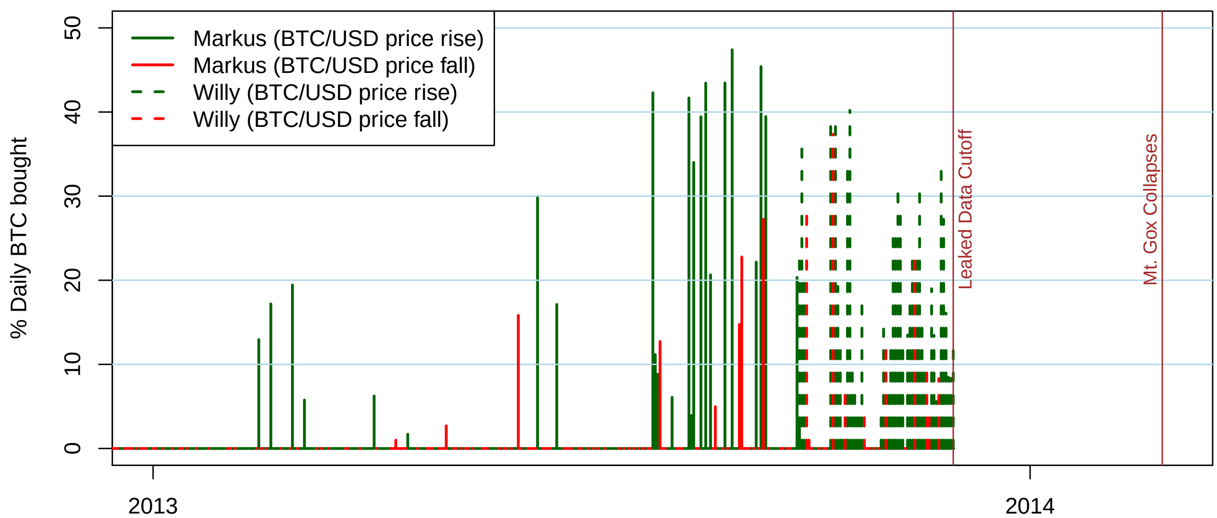

Figure 2: Percentage of total daily trade volume at Mt. Gox when Willy and Markus are active; shaded green if the BTC/USD exchange rate closed higher and red otherwise.

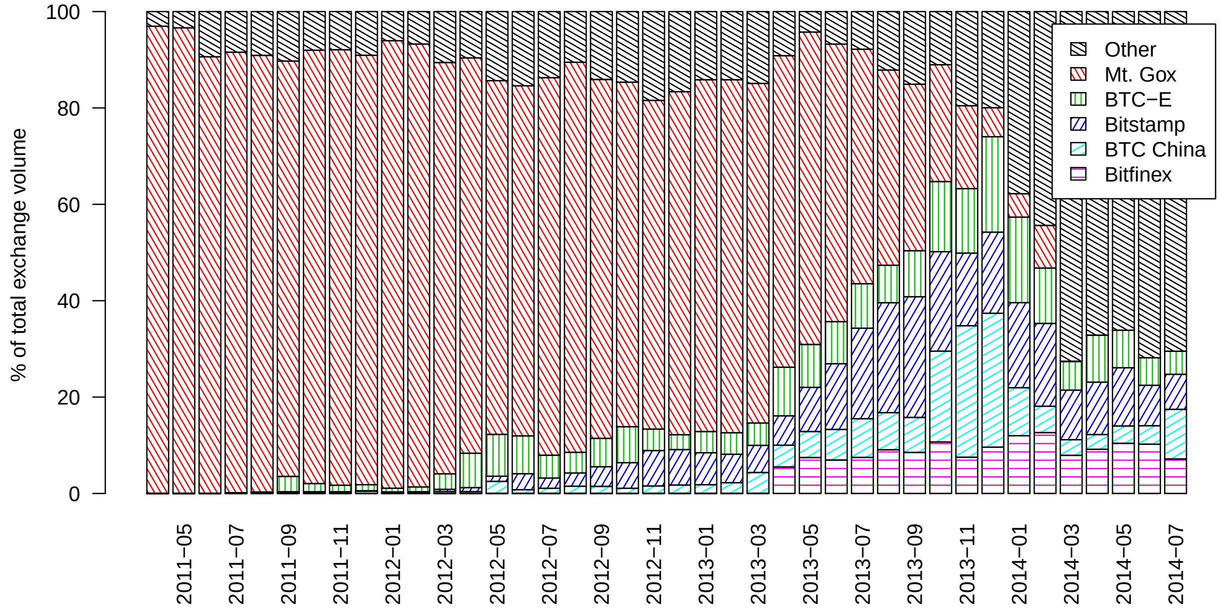

Figure 3: Distribution of market share among Bitcoin currency exchanges by reported trade volume, April 2011 to July 2014. (Source: `bitcoincharts.com`)

List of Tables

Table 1: Daily BTC purchased by Markus and Willy on days they were active.

	Mean	SD	Median	N
Markus:				
BTC purchased	9,302	7,310	5,874	33
% of Mt.Gox daily trade	21		17	
% of total trade at 4 main exchanges	12		10	
Willy:				
BTC purchased	4,962	4,462	3,881	50
% of Mt.Gox daily trade	18		15	
% of total trade at 4 main exchanges	6		5	

Table 2: Comparison of daily BTC volumes on days when suspicious trades occurred and did not.

				Daily BTC Volume		
Buyer	Period	Bot?	Exchange	Mean	Median	N
Markus	3	Active	Mt. Gox	10,056	8,901	17
Everyone	3	Active	Mt. Gox	39,619	42,022	17
Everyone	3	Inactive	Mt. Gox	27,672	17,421	75
Everyone	3	Active	Overall	63,984	67,691	17
Everyone	3	Inactive	Overall	46,962	31,173	75
Willy	4	Active	Mt. Gox	4,962	3,881	50
Everyone	4	Active	Mt. Gox	30,854	25,939	50
Everyone	4	Inactive	Mt. Gox	17,472	10,444	41
Everyone	4	Active	Overall	90,611	82,779	50
Everyone	4	Inactive	Overall	46,263	29,476	41

Table 3: Unauthorized activity and price changes on Mt. Gox

		Days with no bots		Days with bots	
		Days	%	Days	%
Markus	Daily rate decrease	84	44	7	21
	Daily rate increase	109	56	26	79
Willy	Daily rate decrease	9	60	10	20
	Daily rate increase	6	40	40	80
Total	Daily rate decrease	93	45	17	21
	Daily rate increase	115	55	65	79

Table 4: Suspicious trading activity: % of days active during each period

	Period 1 2012-12-01 – 2013-02-28	Period 2 2013-03-01 – 2013-05-31	Period 3 2013-06-01 – 2013-08-31	Period 4 2013-09-01 – 2013-11-30
Markus	3%	5%	19%	9%
Willy	0	0	0	55%
N	90	92	92	91

Table 5: Average daily rate change (in $) and percentage rate change (in parentheses) in USD-BTC exchange rate by period

	Period 1	Period 2	Period 3			Period 4		
			All	Markus active	Markus not active	All	Willy active	Willy not active
Rate change	0.21	1.00	0.16	3.15	-0.51	11.61	21.85	-0.88
Mt.Gox	[1%]	[1.8%]	[0.2%]	[2.9%]	[-0.4%]	[2.6%]	[5%]	[-0.2%]
Rate change	0.23	1.02	0.02	2.35	-0.51	10.99	20.37	-0.45
Bitstamp	[1.1%]	[2.1%]	[0.1%]	[2.3%]	[-0.4%]	[2.6%]	[4.9%]	[-0.05%]
Rate change	.	0.92	0.04	2.14	-0.44	10.75	19.54	0.03
Bitfinex	.	[1.3%]	[0.1%]	[2.2%]	[-0.3%]	[2.7%]	[5%]	[-0.07%]
Rate change	0.22	1.05	-0.1	1.81	-0.53	10.30	19.22	-0.58
BTC-e	[1%]	[2.1%]	[0.01%]	[1.9%]	[-0.4%]	[2.6%]	[4.8%]	[-0.07%]
N	90	92	92	17	75	91	50	41

Table 6: Examining Price Changes Within Mt. Gox and the other Exchanges

Dependent Variable / Independent Variables	Mt.Gox Rate Change	Bitstamp Rate Change	Bitfinex Rate Change	BTC-e Rate Change
Markus	2.79	3.24	2.06	2.37
	(0.72)	(0.96)	(0.31)	(0.71)
Willy	21.65***	20.21***	19.23***	19.04***
	(6.66)	(7.18)	(3.63)	(6.81)
DDoS	-2.38	-1.67	-1.87	-2.01
	(-0.55)	(-0.44)	(-0.26)	(-0.54)
Day After DDoS	-3.50	-3.25	-2.9	-2.68
	(-0.80)	(-0.86)	(-0.41)	(-0.72)
Other Attacks	7.16	5.70	7.35	5.61
	(0.82)	(0.75)	(0.44)	(0.75)
Constant	0.37	0.30	0.45	0.32
	(0.28)	(0.26)	(0.17)	(0.28)
N	365	365	244	365
adj. R^2	0.10	0.12	0.037	0.11

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 7: Examining Percent Price Changes Within Mt. Gox and the other platforms

Independent Variables / Dependent Variable	Mt.Gox % Rate Change	Bitstamp % Rate Change	Bitfinex % Rate Change	BTC-e % Rate Change
Markus	0.0371**	0.0434***	0.0272*	0.0348**
	(3.18)	(3.55)	(1.66)	(2.90)
Willy	0.0433***	0.0423***	0.0469***	0.0413***
	(4.45)	(4.14)	(3.54)	(4.12)
DDoS	-0.0182	-0.00758	-0.00391	-0.00903
	(-1.40)	(-0.55)	(-0.22)	(-0.67)
Day After DDoS	-0.0144	-0.0128	-0.0167	-0.0111
	(-1.10)	(-0.94)	(-0.94)	(-0.83)
Other Attacks	0.0374	0.0234	0.0239	0.0235
	(1.43)	(0.85)	(0.57)	(0.87)
Constant	0.0071	0.0065	0.0032	0.0069
	(1.77)	(1.57)	(0.46)	(1.68)
N	365	365	244	365
adj. R^2	0.075	0.064	0.044	0.054

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Prevalence and Impact of Trading Volume Spikes on Prices in Cryptocurrencies Today

	Days			Rate Change	
Volume	#	%	Currencies	Median	Mean
$\geq 150\%$	$19,212$	8%	304 of 308	1.5%	26.8%
$< 150\%$	$220,988$	92%	–	0%	8.6%
$\geq 200\%$	$14,110$	6%	301 of 308	2%	30.5%
$< 200\%$	$226,090$	94%	–	0%	8.8%

Table 9: Distribution of USD/BTC rates paid by Markus

	$\leq \$0.10$	$> \$0.10$, $\leq \$1$	$> \$1$, $\leq \$20$	$> \$20$, $\leq \$229$	$> \$229$, $\leq \$1K$	$> \$1K$, $\leq \$10K$	$> \$10K$, $\leq \$100K$	$> \$100K$
#	1 050	2 586	6 320	7 009	3 658	4 604	2 311	821
%	3.7%	9.2%	22.3%	24.7%	12.9%	16.2%	8.1%	2.9%

Table 10: Fraudulent transactions initiated by Markus (user ID in bold)

Trade_Id	Date	User_Id	Type	Bitcoins	Money	Money_JPY
1362466099116388	2013-03-05 6:48	238168	buy	0.58932091	22.39419	2094.796
1362466099116388	2013-03-05 6:48	109955	sell	0.58932091	22.39419	2094.796
1362466144485228	2013-03-05 06:49	238168	buy	0.3982007	15.13163	1415.442
1362466144485228	2013-03-05 06:49	132909	sell	0.3982007	15.13163	1415.442
1362466154623847	2013-03-05 06:49	**698630**	buy	1.70382	15.13163	1415.442
1362466154623847	2013-03-05 06:49	96376	sell	1.70382	15.13163	1415.442
1362466154714939	2013-03-05 06:49	**698630**	buy	1	15.13163	1415.442
1362466154714939	2013-03-05 06:49	201601	sell	1	15.13163	1415.442

Table 11: Duplicate Markus Transactions

Trade_Id	Date	User_Id	Type	Bitcoins	Money	Money_JPY
1362466154714939	2013-03-05 06:49	201601	sell	1	15.13163	1415.442
1362466154714939	2013-03-05 06:49	**698630**	buy	**1**	**15.13163**	1415.442
1362466154714939	2013-03-05 06:49	201601	sell	1	38.11000	3564.883
1362466154714939	2013-03-05 06:49	**634**	buy	**1**	**38.11000**	3564.883

Table 12: Summary of Markus transactions

	User ID	# Transactions	Total BTC	Total USD
Manipulated Buy	698630	2966	67 451.61	$1.1M
Manipulated Buy	634	2966	67 451.61	$2.0M
Unchanged Buy	698630	25407	268 446.09	$76.4M
Manipulated Sell	698630	196	5 049.86	$0.2M
Manipulated Sell	634	196	5 049.86	$0.2M
Unchanged Sell	698630	2 927	35 867.18	$4.0M

Table 13: Summary statistics of independent and dependent variables

	Mean	SD	Min	Max
Markus	0.09	0.29	0	1
Willy	0.14	0.34	0	1
DDOS	0.08	0.27	0	1
Day after DDOS	0.08	0.27	0	1
Other Attacks	0.02	0.13	0	1
Mt.Gox daily rate change ($)	3.24	22.39	-139.78	257.5
Bitstamp daily rate change ($)	3.06	19.53	-132.99	190.91
Bitfinex daily rate change ($)[18]	4.25	33.30	-295.97	294
BTC-e daily rate change ($)	2.86	19.28	-134.30	198.14
Mt.Gox daily % rate change	1.4%	6.6%	-28%	49%
Bitstamp daily % rate change	1.5%	6.9%	-49%	40%
Bitfinex daily % rate change[19]	1.4%	8.4%	-37%	59%
BTC-e % daily rate change	1.4%	6.7%	-50%	41%
N	365			

Table 14: Correlation between daily rate changes and the independent variables

	Mt.Gox Rate Change	Bitstamp Rate Change	Bitfinex Rate Change	BTC-e Rate Change
Markus	0.001	0.01	-0.02	0.00009
Willy	0.33	0.35	0.23	0.34
DDoS	-0.06	-0.06	-0.05	-0.06
Day After DDoS	-0.07	-0.07	-0.05	-0.06
Other Attacks	0.02	0.02	0.013	0.02
N	365	365	244	365

Table 15: Correlation between daily percent rate changes and the independent variables

	Mt.Gox % Rate Change	Bitstamp % Rate Change	Bitfinex % Rate Change	BTC-e % Rate Change
Markus	0.14	0.16	0.07	0.13
Willy	0.21	0.2	0.22	0.2
DDoS	-0.1	-0.05	-0.05	-0.06
Day After DDoS	-0.09	-0.06	-0.08	-0.06
Other Attacks	0.07	0.04	0.02	0.04
N	365	365	365	365

Table 16: Correlation between independent variables

	Markus	Willy	DDoS	Day After DDoS	Other Attacks
Markus	1				
Willy	-0.1	1			
DDoS	0.05	-0.06	1		
Day After DDoS	0.05	-0.06	0.33	1	
Other Attacks	0.03	-0.05	-0.04	0.04	1
N	365				

Table 17: Suspicious trading activity and price changes on Bitstamp

		Days with no STA		Days with STA	
		days	%	Days	%
Markus	Daily rate decrease	88	45	6	18
	Daily rate increase	105	55	27	82
Willy	Daily rate decrease	6	40	9	18
	Daily rate increase	9	60	41	82
Total	Daily rate decrease	94	45	15	18
	Daily rate increase	114	55	67	82

Table 18: Willy: Volume activity (period 4)

	mean	median	N
Volume bought by Willy (Mt. Gox)	4,962	3,881	50
Total BTC volume on Mt. Gox (Willy active)	30,854	25,939	50
Total BTC volume on Mt. Gox (Willy inactive)	17,472	10,444	41
Total BTC volume on Bitstamp (Willy active)	26,084	23,684	50
Total BTC volume on Bitstamp (Willy inactive)	14,793	10,505	41
Total BTC volume on Bitfinex (Willy active)	12,981	11,756	50
Total BTC volume on Bitfinex (Willy inactive)	6,467	3,829	41
Total BTC volume on BTC-e (Willy active)	20,691	18,661	50
Total BTC volume on BTC-e (Willy inactive)	7,529	3,737	41
Total BTC volume (Willy active)	90,611	82,779	50
Total BTC volume (Willy inactive)	46,263	29,476	41

Table 19: Markus: Volume activity (period 3)

	mean	median	N
Volume bought by Markus (Mt. Gox)	10,056	8,901	17
Total BTC volume on Mt.Gox (Markus active)	39,619	42,022	17
Total BTC volume on Mt.Gox (Markus inactive)	27,672	17,421	75
Total BTC volume on Bitstamp (Markus active)	13,547	12,840	17
Total BTC volume on Bitstamp (Markus inactive)	10,299	8,850	75
Total BTC volume on Bitfinex (Markus active)	5,976	5,622	17
Total BTC volume on Bitfinex (Markus inactive)	4,331	3,197	75
Total BTC volume on BTC-e (Markus active)	4,840	4,699	17
Total BTC volume on BTC-e (Markus inactive)	4,660	3,589	75
Total BTC volume (Markus active)	63,984	67,691	17
Total BTC volume (Markus inactive)	46,962	31,173	75