

The Law of One Bitcoin Price?

Alexander Kroeger and Asani Sarkar¹

Abstract

Bitcoin, a digital currency, constitutes a textbook example where the law of one price should be satisfied as, unlike asset pairs previously studied in the literature, it is a fully fungible asset with identical payoffs. Despite this, we show the existence of persistent, statistically significant differences between US dollar-denominated bitcoin prices in multiple bitcoin exchanges. We argue that the price difference has two components: One, microstructure frictions such as the bid-ask spread, order book depth and volatility and negatively related to volume; and market segmentation. Two, an “illegality premium” due to the use of bitcoin to avoid foreign exchange restrictions, money laundering, etc. We find evidence that microstructure frictions explain part, but not all, of the price differences. For example, the absolute values of price differences are positively related to the bid-ask spread, order book depth and volatility and negatively related to volume. Price differences are also higher on exchanges with smaller trade sizes, consistent with clientele effects from greater institutional trading. Moreover, impulse responses indicate that shocks to illiquidity and volatility have more persistent effects on absolute price differences between exchange pairs with more retail trading and greater counterparty risk. Finally, the speed of arbitrage and the amount of price discovery is related to the arbitrage frictions. Thus, limits to arbitrage remain relevant even in the context of a homogeneous asset class where many of the frictions in more traditional asset markets are absent. In ongoing work, we quantify this illegality premium using a number of innovative measures of illegality.

Keywords: bitcoin exchanges, limits to arbitrage, digital currency

JEL classification: G12, G14

The views expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or of the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

¹ Federal Reserve Bank of New York; asani.sarkar@ny.frb.org

1 Introduction

The law of one price states that assets with identical payoffs should trade at the same price but, on occasion, this law appears to be violated. Examples of such violations are “Siamese-twin” stocks with almost-identical dividend streams that trade at different prices (Jong et al. (2009), Rosenthal and Young (1990) and Debora and Froot (1999)), parent and subsidiary company stocks trading at prices such that the parent company value is negative (Mitchell et al. (2002), Lamont and Thaler (2003)) and the off-the-run minus on-the-run Treasury bond spread (Amihud and Mendelson (1991), Warga (1992) and Krishnamurthy (2009)). Nevertheless, this evidence exists for only a limited set of asset pairs since those with closely related payoffs are hard to come by (Gromb and Vayanos (2010)). Additional violations of the law of one price come from cross-listed stocks, such as American Depositary Receipts (ADR), where the stock of the same company trading in geographically-dispersed exchanges trade at different prices (Gagnon and Karolyi (2010)). In all these cases, although the payoffs are close, they are not identical. For example, even ADRs and the foreign shares they represent are not fully fungible (Gagnon and Karolyi (2010)).

In this paper, we study price differences of the virtual currency bitcoin trading on 6 exchanges constituting 15 exchange pairs where the price is quoted in US dollars (USD). Unlike asset pairs previously studied in the literature, bitcoin in different USD exchanges is by design the same, fully fungible, asset with identical payoffs and no currency risk. As such, these exchanges constitute ideal, textbook examples where the law of one price should be satisfied. Our results, however, show the existence of persistent, statistically significant differences between bitcoin prices quoted in multiple USD exchanges. These price differences constitute two groups. For exchange pairs involved BTC-e, bitcoin *consistently* trades at substantially lower prices on BTC-e compared to the other exchanges, and the average price difference is statistically significantly different from zero. For the remaining exchange pairs, price differences are typically more *episodic* and smaller in magnitude, and they switch signs.

For example, higher transfer costs and lower liquidity are likely to result in larger trading frictions and thus higher price discrepancies. Consistent with this idea, we find that the absolute values of price differences are positively related to the bid-ask spread, order book depth and volatility and negatively related to volume. We also account for explicit fees, in particular bitcoin network trading fees that users specify in order to incentivize so-called “miners” who create bitcoin.¹ However, even after accounting for the implicit and explicit trading costs, we find that the average absolute price difference remains statistically significant for exchange pairs involved BTC-e.

¹See Section 2.1 for a more detailed description.

Since trading frictions do not fully account for price differences in the exchange pairs involving BTC-e, we next examine market segmentation and institutional factors as determinants of bitcoin arbitrage profits. Segmentation may arise from clientele effects with the more active exchanges hosting more institutional traders, as inferred from differential trade sizes. We define dummy variables *bothretail* and *oneretail* that indicate whether both or one of the exchange pairs, respectively, have small trade sizes and find that absolute price differences are positively related to the incidence of retail trading.² We account for institutional factors through exchange-pair fixed effects. These fixed effects are generally statistically significant, implying that institutional features matter. Other factors, such as anonymity and exchange failure risk are also important, although we cannot quantify them. In particular, BTC-e offers enhanced anonymity relative to other exchanges, making it potentially attractive to users wishing to convert illicitly-obtained bitcoin into fiat currency, thus creating increased selling pressure (and hence lower prices) compared to other exchanges. Traders may also perceive BTC-e as carrying a greater risk of loss of customer accounts due to the opacity of their operations, entailing a risk premium on an arbitrage trade.

Do trading frictions account for the persistence of price differences? We estimate impulse responses from a Vector Autoregression (VAR) of price differences and search frictions. We find that shocks to illiquidity and volatility have persistent effects on price differences between BTC-e related exchange pairs, lasting 10 days or more. In contrast, the effect of shocks on price differences are more attenuated, lasting 2 to 3 days, for exchange pairs that do not involve BTC-e. Overall, both the magnitude and persistence of price differences are related to search frictions.

How do arbitrage frictions affect price efficiency dynamics and price discovery? We estimate a Vector Error Correction Model (VECM) and find that prices of all exchange pairs are cointegrated of order one, indicating the existence of a long-run equilibrium relation between prices of every exchange pair. However, there are substantial differences between exchanges in the speed of adjustment of prices to deviations from the long-run equilibrium relation. Specifically, the half-life of a price deviation (i.e. the number of days required to eliminate half of the price deviation) is two to three times higher for BTC-e, as compared to other exchanges. Moreover, the information share of BTC-e is small compared to other exchanges, indicating that little price discovery occurs on this exchange for the period analyzed. By contrast, a substantial part of price discovery occurs on Bitfinex, relative to other exchanges. These results imply that arbitrage frictions impede the speed with which prices revert to equilibrium and the amount of price discovery.

² “Retail exchanges” are defined as having a 75th percentile of trade size that is below the 75th percentile of all trade sizes for the past 30 days.

We contribute to the literature by showing that search frictions, market segmentation and institutional features all contribute to limiting arbitrage between bitcoin exchanges in the cross-section and in the time series. We demonstrate how frictions affect arbitrage profits across exchange pairs. Further, we relate the persistence of arbitrage profits to search frictions, and show how the dynamics also differ across exchange pairs. Finally, we relate arbitrage frictions to price discovery. Thus, limits to arbitrage remain relevant for price efficiency even in the context of a homogeneous asset class where many of the frictions in more traditional asset markets are absent.

The article is organized as follows. Section 2 provides an overview of Bitcoin and the exchange market. Section 3 provides more information on the sources of data used to carry out the analysis. Section 4 documents the price relationships between the major USD–BTC exchanges and describes the various arbitrage frictions that could lead to price discrepancies. Section 5 presents regression results to quantify the effects of various frictions on price disparities. Section 6 analyzes some of the dynamic considerations for arbitrage. Section 7 concludes.

2 How bitcoin is owned, used, priced and regulated

Bitcoin is a digital asset and payment system introduced in 2009 that has attracted a great deal of attention due to its innovative design. Over the last seven years, bitcoin has grown to become a major medium of exchange with an estimated 100,000 merchants worldwide accepting bitcoin.³ In this section, we describe the use of bitcoin as a decentralized payments system (section 2.1) and how bitcoins are traded over online exchanges (section 2.2). Regulatory issues related to exchanges are discussed in section 2.3.

2.1 The bitcoin payment system

Ownership of the units of bitcoin is established through claims to the current state of the public transaction history known as the “blockchain.” Users of bitcoin are identified by one or more bitcoin “addresses.” These addresses are known to the public and are often managed by client software known as a “wallet.” In order to claim ownership of bitcoin, that is, to send bitcoin from a particular address, one must possess the “private key” associated with that address. As the name suggests, the private key is not known to the public. It can be used to create a verifiable digital signature for

³60,000 merchants were accepting bitcoin through Bitpay alone (<https://blog.bitpay.com/bitcoin-a-new-global-economy/>).

sending bitcoins from one's address, thereby demonstrating that the transaction is valid (as long as the private key has not been compromised).

One concern with any decentralized payments system is that the currency units promised in a transactions have already been spent, i.e. how does the payee know that the payer has not already pledged the bitcoins in their wallet to another user. This is referred to as the “double spend” problem. In the context of the Bitcoin network, a user could sign multiple transactions referring to the same entry on the blockchain. The process of adding transactions to the blockchain and verifying that these transactions do not involve previously spent bitcoins is called “mining.” Miners add transactions to the blockchain in groups called “blocks.” In order to provide an economic incentive for mining and to create new bitcoins, miners are awarded a fixed amount of bitcoins per block⁴ as well as any transaction fees specified by the senders. The rate at which blocks are mined is automatically adjusted by the network such that a block is mined roughly every 10 minutes, thus controlling the rate at which new bitcoins are created.

The size of these blocks is currently capped at 1 megabyte, although there is ongoing debate about whether or not to increase the block size limit, since the current 1 megabyte is quickly becoming inadequate for the number of transactions occurring over the network (see Figure A1 in the appendix for a plot of the blocksize). As of this writing, the block size remains capped at 1 MB, although software changes have been proposed to increase the limit and thus the capacity of the network.⁵ As blocks push against the blocksize limits, users experience higher transaction fees and/or time delays in having transactions added to the blockchain.

One of the novel features of the bitcoin protocol is that it lacks any central authority. The protocol as just described is determined by a set of code that is valid insofar as it is adopted by miners and users. In this way, bitcoin contrasts with fiat currencies, such as the dollar and the euro, which are issued and regulated by a central authority (such as a central bank) and constitute legal claims on their issuers. For example, bank deposits are the claims of the assets of banks and Federal Reserve notes (such as dollar bills) are technically claims on the assets of the Federal Reserve.

2.2 The bitcoin exchange market

Price volatility has been a feature of bitcoin since active trading between the U.S. dollar and bitcoin began (Figure 1). These prices are expressed over online exchanges which allow customers to ex-

⁴The reward per block is halved roughly every four years. As of August 2016, 12.5 bitcoins are created for each new block of transactions.

⁵A group of Bitcoin developers have proposed a revised protocol that incorporates an increase in the blocksize limit to 2 MB called **Bitcoin Classic**. This protocol is competing with **Bitcoin Core** which aims to keep the blocksize limit at 1 MB.

change bitcoins for fiat currency (or in some cases, other digital currencies). Thus, these exchanges are the primary mode of obtaining bitcoin. They act either as brokers, providing a platform over which buyers and sellers can meet for a fee, or as dealers that profit from bid-ask spreads.

Online exchanges are currently an important source of price discovery for bitcoin—they are the mechanism through which bitcoin’s value is expressed, allowing it to serve as a medium of exchange. Unlike major fiat currencies, bitcoin does not (at least not yet) serve as a widely quoted unit of account in and of itself. As [Ali et al. \(2014\)](#) point out, few prices are actually negotiated in bitcoin. Rather, in situations where bitcoin is accepted as payment, price setting typically occurs in terms of fiat currency—the unit of account—and bitcoin prices are quoted in terms of a conversion of the fiat price at current exchange rates.⁶ Relatedly, many companies that accept bitcoin, such as Dell and Microsoft, never actually receive bitcoin. Rather, they use third parties such as BitPay and Coinbase to process bitcoin payments from their customers and then forward dollars (or another fiat currency) to the merchant.

Historically, the bitcoin exchange market has been dominated by a few large players. For example, prior to its demise February 2014, the exchange Mt.Gox handled the vast majority of dollar to bitcoin transactions. Figure 2 shows the volume of bitcoin traded over exchanges over time, broken out by individual exchanges. After the collapse of Mt. Gox in 2014, three exchanges largely filled the void in the dollar–bitcoin market: Bitfinex, Bitstamp, and BTC-e. Recently, itBit and Coinbase have also emerged as significant exchanges.

2.3 Regulation of bitcoin exchanges

Due to its anonymous nature (i.e. a public address that is not associated with any identity), bitcoin has become a widely used medium of exchange for the so-called “Dark Web.” It was infamously used on the online marketplace for Silk Road, nicknamed the “eBay of drugs,” which was shut down by the FBI in 2013 ([Caffyn \(2015\)](#)). It has also been used a tool for hackers to collect ransoms from their targets.⁷

Given that bitcoin exchanges involve transfers of (potentially large amounts) of value, they fall under anti-money laundering (AML) and know-your-customer (KYC) regulations designed to hamper criminal activity by de-anonymizing financial transactions. However, the extent to which these rules apply depend on the particular jurisdiction of the exchange’s operations and of the customers with whom the exchange operates. [Pieters and Vivanco \(2016\)](#) explores how varying levels of com-

⁶For instance, Overstock.com, an early adopter of bitcoin, provides an option at checkout to allow the customer to pay the USD equivalent in bitcoin. Refunds are also made in the USD value of the item, not the amount of bitcoin paid.

⁷For an example, see the [LA Times](#)

pliance (and therefore varying levels of anonymity) affect prices across exchanges, a topic that we explore in Section 4.

3 Data

Bitcoin lends itself to economic study due to its great quantity of publicly available information. By design, the blockchain is a public record of all transactions ever sent over the bitcoin network. Many organizations, such as exchanges, involved in the Bitcoin ecosystem have adopted the ethos of transparency.

Our primary source of exchange data is bitcoincharts.com. Many major bitcoin exchanges voluntarily maintain an API feed of their trades and order book. Bitcoincharts.com aggregates the entire trade histories of exchanges that conform to its formatting standards and makes them available for download. Many exchanges participate, which makes it a rich source of data, but it is not comprehensive. For example, notable absences from Bitcoincharts are Gemini, OkCoin and LakeBTC. According to another bitcoin data aggregation site—bitcoinity.org—that does not have trade histories but does track total volume, the bitcoincharts.com data covers about 70% of the BTC-USD market from the beginning of 2015 to August 2016. From bitcoinity.org, we also pull summarized order book information—namely bid-ask spreads as well as the sum of orders at a daily frequency for major exchanges.

The available data by exchange is shown in the appendix in Tables B1 and B2. To easily compare prices, we confine our sample only to exchanges where the price is quoted in USD. An observation in the transaction data consists of a UTC (coordinated universal time) timestamp, the amount of bitcoin traded, and the USD price at which the trade occurred. Virtually all exchanges operate continuously with the exception of operational outages, and so business hours are not observed in the data.

For the analysis in this paper, six exchanges are considered—Bitfinex, Bitstamp, BTC-e, Coinbase, itBit, and Kraken—which yields 15 unique pairs. Mt. Gox is omitted due to the lack of daily information on its liquidity. Trades are time stamped down to the second defined by coordinated universal time (UTC) for consistency. For the analysis we aggregate price the price difference series to the daily level. The prices differences are first calculated by taking the difference of volume weighted prices across exchanges over five minute intervals. We then take a volume weighted average of these five minute price differences at the daily level in order to obtain a daily price difference series.

4 Deviations from the law of one price

In theory, bitcoin prices should be identical across exchanges. Bitcoin is a truly homogeneous asset and highly transferable. Bitcoins are perfectly fungible—one bitcoin is substitutable for any other. Furthermore, bitcoin ownership can be transferred in a short amount of time at low cost.⁸ Therefore, an arbitrageur could in theory safely profit by buying bitcoin on an exchange where it is less expensive and then selling it or establishing a short position on an exchange where it is more expensive. This mechanism should enforce the Law of One Price, whereby the price of bitcoin should be the same regardless of where it is purchased. We estimate price differences across exchange pairs in section 4.1 and show that they are significantly different from zero and persistently so. In section 4.3, we describe in greater detail the strategy an arbitrageur might employ to take advantage of a difference between bitcoin prices in two exchanges. Section 4.2 discusses two sources of price differences: risk premia resulting from illegal activities and various limits to arbitrage impeding the execution of this strategy.

4.1 Price Differences between Bitcoin Exchange Pairs

Figures 3 and 4 show plots of daily time series of price differences between exchanges, normalized by the average price between the two exchanges, and expressed as a percent. The price difference is capped in the range -10% to $+10\%$ in order to show comparisons of similar magnitude, but the maximum absolute difference can be higher, as indicated below each chart. The sample period is based on data availability which differs by exchange pair, as reported in the appendix table B2. The magnitude of the normalized price difference is shown as a blue area plot. Each plot has an orange line indicating the average unsigned price difference, and a red line showing the average absolute price difference between the two exchanges.

Figure 3 shows the price difference as the price on exchange 1 (one of the four major bitcoin exchanges Bitfinex, Bitstamp, Coinbase or ItBit) minus the bitcoin price on BTC-e. For example, a value of 0.10 on the graph labeled “Bitfinex - Btce” indicates that bitcoin on Bitfinex traded at a 10% premium on that day, relative to bitcoin on BTC-e. Bitcoin on BTC-e consistently trades at a 1-2% discount or more relative to the other exchanges. Both the average signed difference and the average absolute difference is roughly 2% for each exchange pair, indicating that the sign of the price difference rarely flips. However, the maximum absolute price difference ranges between 17% for the Coinbase-BTC-e exchange pair and as high as 41% for the Bitstamp-BTC-e pair. The p-

⁸Transaction fees related to transferring bitcoin have been rising in recent months due to block size constraints. For our purposes, especially for large bitcoin transactions, the fees can be considered negligible.

values shown below the graphs reject the null hypothesis that the average unsigned price difference equals zero at less than a 1% level of significance for every exchange pair. SHOULD WE ALSO SHOW A TEST OF THE ABSOLUTE PRICE DIFFERENCE?

Figure 4 shows the bitcoin price on exchange 1 (one of Bitstamp, Coinbase, itBit, and Kraken exchanges) minus the bitcoin price on Bitfinex. Compared to the price differences relative to BTC-e, these exchange pairs tend to have smaller, less persistent price differences. Moreover, the price differences switch signs, indicating that bitcoin prices on Bitfinex are not persistently at a discount or a premium relative to other exchanges. Therefore, while the mean absolute price difference is 1%, the mean unsigned price difference is less than 0.5%, although the latter is statistically different from zero in some cases. Nevertheless, these exchanges exhibit large price discrepancies for multiple days, as indicated by the maximum absolute price deviation that range between 16% and 38%.

Overall, we find large and persistent bitcoin price differences between BTC-e and other exchanges. Moreover, these price differences are typically one-sided in that bitcoin generally trades at a discount on BTC-e relative to other exchanges. In contrast, price differences between Bitfinex and other exchanges are smaller and less persistent, and the sign of the difference flips often.

4.2 Sources of Price Deviations on Bitcoin Exchanges

Although the persistent discount of BTC-e prices has been noticed by users, the price difference remained in force at the end of our sample.⁹ Referring to Figures 3 and 4, there appear to be two separate sources of deviations from price parity. First, there are short-run deviations that dissipate relatively quickly. These are likely caused by “search frictions” that delay trades because of the need to find a trading counterparty (Duffie (2010)). In the context of bitcoin exchanges, where trading is frequent, search frictions are likely to result in price differences that revert relatively quickly. Moreover, these deviations could be two-sided, unless the magnitude of frictions is consistently greater on one exchange relative to another. The second component is a longer-run deviation from the law of one price, particularly observable in the price differences with respect to BTC-e, where bitcoin consistently trades at a discount relative to prices on other exchanges. The main source of these long-lasting price discounts is likely to be a risk premium related to illegal activities, as trades attracted by lower levels of compliance with KYC/AML regulations push down prices on BTC-e relative to other exchanges. An additional source of risk premium arises from the requirement to compensate traders for the risk of exchange failure.

⁹See this [Reddit post](#) or this [Stack Exchange post](#) for examples of users discussing the arbitrage.

In this sub-section, we discuss risk premia arising from the likelihood of trading related to illegal activities on bitcoin exchanges, and the possibility of exchange failure. In sub-section 4.3, we discuss the search frictions that might prevent short-run arbitrage from closing the price differences between bitcoin exchanges.

Bitcoin traders engaged in illegal activities would prefer high levels of anonymity. As described in section 2.3, KYC compliance refers to the requirement for financial institutions to be able identify their customers, should the need arise to examine financial transactions to track illicit activity. These regulations reduce the anonymity of clients by design. Many exchanges exercise KYC compliance to some degree in order to be compliant with the regulations of their respective jurisdiction. An exception is BTC-e which lacks KYC compliance, as it requires very little personal information in order to transact over its platform (Pieters and Vivanco (2016)).¹⁰ URL IN FOOTNOTE IS NOT CLICKABLE. This enhanced anonymity could result in increased selling pressure (and thus lower prices) if users convert their illicitly-obtained bitcoin to fiat currency on BTC-e, especially since no other major USD exchange is sufficiently anonymous to serve this purpose.

Differences in the risk of exchange failure are also relevant for understanding persistent bitcoin price discounts. Moore and Christin (2013) found that 18 of 40 early bitcoin exchanges studied by them closed, often with customer balances wiped out. Since bitcoin's sole proof of ownership is the possession of a public address' private key, hacking is a concern. If hackers discover the private keys used for the exchange's balances, they can transfer bitcoin to other accounts controlled solely by them. The most notable bitcoin exchange failure has been that of Mt. Gox, an exchange that once commanded nearly the entirety of the bitcoin market (Figure 2), prior to losing roughly \$460 million of the value of its users' bitcoins to hackers and failing in February 2014 (McMillan (2014)). In January 2015, hackers successfully transferred funds out of Bitstamp accounts, although Bitstamp was able to cap its losses at \$5.1 million and continue to operate and honor all customer claims (Russell (2015)). In August 2016, the Hong-Kong based exchange Bitfinex, currently the largest bitcoin-USD exchange, lost roughly \$70 million of its clients' bitcoin. Bitfinex announced that it would continue operating, but would reduce all customer account balances by 36% (Tepper (2016)). CAN WE SHOW THAT BITCOIN TRADES AT A DISCOUNT FOLLOWING HACKS IN THESE CASES? LOOK AT VOLUME TOO. IF NOT, THEN THE EVIDENCE DOES NOT SUPPORT WHAT WE'RE SAYING HERE. Differences in the risk of losing account balances in hacks may lead to difference arbitrage premiums across different exchange pairs to compensate arbitrageurs for interacting with risky counterparties. Compounding BTC-e's failure risk is the fact that, although it has been operating longer than any other major USD-bitcoin exchange, little is

¹⁰For example, see this June 2016 post or this May 2015 review.

known about its owners or even the location of its activities, and so there is likely limited legal recourse for investors.¹¹

4.3 Bitcoin Arbitrage Strategy

Large price deviations between bitcoin exchange prices suggest frictions to acting on the potential arbitrage opportunity (Shleifer and Vishny (1997)). In order to exploit bitcoin price differences across exchanges, an arbitrageur could execute a simple trade: buy bitcoin on the exchange where it is relatively cheap and sell it where it is relatively expensive. The mechanics of this trade is shown diagrammatically in Figure 5 using BTC-e and Bitfinex as an example where bitcoin is priced relatively lower on BTC-e and relatively higher on Bitfinex (often the case in the data, per Figure 3). The sequence of steps in the arbitrage trade is as follows.

- To fund the trade, the arbitrageur transfers dollars into an account with BTC-e, which entails a fee and a delay of several business days.
- The arbitrageur purchases bitcoin on BTC-e at the “ask” price, paying a trading fee.
- He or she either (a) shorts bitcoin at the “bid” price on Bitfinex and transfers the bitcoin purchased on BTC-e to cover the short position, paying a margin funding fee or (b) transfers the bitcoin from BTC-e to an account on Bitfinex and sells it at the “bid” price after the transfer, paying a trading fee.
- Finally, the arbitrageur transfers dollars out of the Bitfinex account to realize profits on the trade, incurring fees and experiencing time delays in the wire transfer.

Delays caused by search frictions are expected to be correlated with higher trading fees, bid-ask spreads and price volatility, and with lower order-book depth and trading volume. Figure 5 illustrates that each step of the arbitrage trade involves frictions in the form of various fees and time delays. Fees arising from trading bitcoins stem from the both the exchanges and the bitcoin network. Network transaction fees are used in order to incentivize miners to include the transaction in their blocks (see section 2.1), thus leading to faster confirmation by addition to the blockchain. As the size of blocks has approached the 1 MB limit imposed by the protocol, the need to provide fees for faster confirmation times has become more acute. Figure 6 shows how the average transaction fee per transaction has evolved over time. An arbitrageur would incur the transaction fee when

¹¹Various guesses for the location of BTC-e’s operations proposed by one source are Bulgaria, Cyprus, and Russia—see [bitreview](#).

depositing bitcoin in the higher priced exchange in order to sell it or to cover a short trade. As these fees are relatively small in magnitude (a few cents per transaction), it is not clear a priori to what extent they contribute to the price differences.

In addition, exchanges typically charge a trading fee as a percentage of each transaction. We do not have a time series of the trade-specific fees. However, Table 1 summarizes typical fees as of August 2016 for the major USD–bitcoin exchanges, obtained from the exchange websites. Note that trading fees are generally highest for BTC-e, which also has the highest price differences with respect to other exchanges. For these fees, it is common for exchanges to distinguish between “taker” and “maker.” The side that crosses the bid-ask spread is considered the “taker,” whereas the side whose order is filled is considered the “maker.” To encourage participants to maintain market depth in the order book, “maker” trades generally entail smaller fees. In order for a trader to pursue an active arbitrage strategy, they would need to cross the bid-ask spread (a “taker” trade) to ensure that their order is filled. Alternatively, the trader can pursue a passive arbitrage strategy in which they place an order on the other side of the bid-ask spread in order to receive a more favorable price as well as pay a lower fee, but this opens the trader up to the risk that the order will not be filled in a timely manner, and the price could move in the opposite direction before the order is filled. Exchanges may charge additional fees for deposits into and withdrawals from a USD account with the exchange (for example, through a wire transfer).

Market depth is also pertinent to the arbitrageur. The amount of orders available near the inside quotes must be of sufficient quantity to make it worthwhile for the arbitrageur to exploit the price difference without incurring a large and adverse price impact. For the “passive” strategy described above, volume is also relevant so that submitted orders are filled in a timely manner.

Further delays in executing bitcoin arbitrage trades occur due to delays in depositing dollars into an exchange account. For example, on BTC-e, where an arbitrageur would often want to purchase bitcoin with dollars to execute an arbitrage trade, deposits of U.S. dollars via wire take five to ten days to complete.¹² Over this time frame, the expected arbitrage profits become quite uncertain. A measure of uncertainty is price volatility since increased volatility means that the price differences can shrink or even revert before the trade can be fully executed. Indeed, bitcoin prices are highly volatile: the intraday standard deviation of the bitcoin price index often exceeds the average bitcoin price difference between BTC-e and Bitfinex and the two series appear correlated (Figure 7).

Another delay in executing arbitrage trades, though a shorter one, occurs due to the time needed to transfer bitcoin from one exchange to another. A bitcoin transfer on the blockchain takes roughly 10 minutes to settle (to be added to a block). Often exchanges will not consider a transfer valid

¹²See BTC-e’s [website](#)

unless it is several blocks deep in the blockchain.¹³ For instance, Bitfinex and Bitstamp each require three network confirmations, or roughly 30 minutes. This delay is potentially avoidable by short selling bitcoin on the more expensive exchange, but shorting is a service currently only offered by Bitfinex and Kraken, and it entails additional fees.

It is widely acknowledged that arbitrage is mainly carried out by specialized, highly sophisticated traders (Shleifer and Vishny (1997)). There is evidence that different bitcoin exchanges serve different types of clients. Figure 8 shows the distribution of transactions sizes of bitcoin-USD trades (denominated in bitcoins) across different exchanges from January to August of 2016. Bitfinex and itBit facilitate larger trades than other Bitcoin-USD exchanges, suggesting that they may be serving larger, “institutional” clients (such as traders and market makers). Other exchanges typically show smaller transaction sizes, which may indicate that they are primarily serving retail users of bitcoin. The types of traders that generate large orders may be more likely to engage in arbitrage activities, thus making the prices on exchanges that cater to these clients more integrated. Market segmentation is also created by differences in exchange governance rules. In particular, BTC-e does not accept wire transfers from US citizens or US banks.¹⁴ While alternative methods of funding an account with USD exist, this could impede many US-based institutional traders from taking advantage of BTC-e’s lower prices.

In summary, our discussion indicates two likely sources of substantial and persistent bitcoin price differences across exchanges. One source arises from risk premia related to poor governance (attracting illegally motivated trades) and the probability of exchange failure. These factors likely results in one-sided price deviations (e.g. BTC-e bitcoin prices being consistently lower than those on other exchanges). Second, fees and trading time delays result in frictions to arbitraging price differences in the short-run. Next, we develop empirical measures for risk premia and arbitrage frictions, and explain price differences by these measures.

5 Determinants of exchange price deviations

Arbitrageurs looking to exploit differences in bitcoin prices are subject to risk from illegality-motivated trading and exchange failure, and from frictions that prevent arbitrage. Arbitrage fric-

¹³Transactions that are “deeper” in the blockchain can be considered more secure, since it reduces the probability of a successful “double spend.” For instance, a miner actively trying to double spend would send bitcoins to the recipient, wait for that transaction to be mined into a block, and then compete against all other miners to create blockchain history without the transaction taking place. The likelihood of success for this strategy decreases as the number of blocks that need to be replace increases.

¹⁴This could be related to BTC-e’s lack of KYC compliance, as discussed in section 2.3. See BTC-e’s [terms and conditions](#).

tions are likely increasing in implicit and explicit transactions costs and the degree of market segmentation while risk premia is created by exchange failure and governance risk. In section 5.1, we examine the relation of price differences to market liquidity, trading fees and price uncertainty. In section 5.2, we examine how impediments to arbitrage relate to market segmentation between retail and institutional traders. In section 5.3, we examine exchange-specific fixed effects to identify the component of price differences related to risk premia.

5.1 Search frictions: illiquidity and volatility

To study the effects of transactions costs and price uncertainty on bitcoin price differences between exchanges, we estimate the following panel regression model using OLS with standard errors clustered by exchange pair.

$$\begin{aligned} PriceDif_{i,t} = & \beta_0 + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_t \\ & + \beta_5 AvgPriceSD_{i,t} + \epsilon_{i,t} \end{aligned} \quad (1)$$

where $PriceDif_{it}$ ¹⁵ represents the absolute value of the bitcoin price difference for an exchange pair i on day t , divided by the the average of the bitcoin prices on the two exchanges:

$$PriceDif_{i,t} = \left| \frac{P_t^1 - P_t^2}{(P_t^1 + P_t^2)/2} \right| \quad (2)$$

All regressors, with the exception of the average price standard deviation $AvgPriceSD$, are log transformed. Further, with the exception of $LnAvgNetworkFee_t$, all of the regressors represent an average of the values for the two exchanges in pair i . $LnAvgBASpread_{it}^1$ is the proportional bid-ask spread (i.e. the bid-ask spread divided by the mid-price). Since the arbitrageur must buy bitcoin in one exchange at the higher “ask” price and sell it at the other exchange at the lower “bid” price, we expect the price difference to be higher when the average bid-ask spread on the two exchanges is higher. $LnAvgOB_{it}^1$ is the daily sum of orders in the order book within 1% of the average daily price, divided by the daily volume. This variable is a measure of the depth of the order book close to the inside quotes, and we expect the price difference to decrease when the order book is deeper. $LnAvgUSDVol_{it}$ is the trading volume in millions of USD, and it is also expected to have a negative relation to the price difference. $LnAvgNetworkFee_t$ is the average transaction fee (in USD) for all transactions added to the blockchain on day t and is common to all exchange pairs. Finally,

¹⁵Note that the price difference series is calculated at 5 minute intervals and then aggregated to the daily level. For a more complete description as to how this is calculated, see Section 3.

$AvgPriceSD_{i,t}$ is the intraday price standard deviation normalized by the price. Increased price volatility creates uncertainty as to the duration of the arbitrage opportunity and thus is expected to increase price differences. $\epsilon_{i,t}$ is an error term.

NEW TABLE 2 WITH DESCRIPTIVE STATISTICS FOR RIGHT HAND SIDE VARIABLES IN REGRESSION, BY EXCHANGE PAIR. COULD BE FORMATTED AS IN TABLE 5.

The results are presented in Table 2. In column (1) of the table, we show results when only the market liquidity measures are included. We find that a higher average spread between two exchanges is associated with a higher absolute price difference, as expected, and this effect is statistically significant. Further, the cumulated orders within 1% of the current price is significantly negatively related to the price difference. In other words, increases in the depth of the order book are associated with decreased price differences. The coefficient on volume is not statistically significant. The network fee is positively related to the price difference but the coefficient is not significant. The lack of significance may be due to the omission of transactions fees and fees to withdraw/deposit USD via wire transfers or other means that may also be material, but are not accounted for in the regression due to the lack of reliable time series information on these factors.

Column (2) of Table 2 shows that the coefficient on the standard deviation is positive and highly significant, suggesting that higher volatility is associated with higher price differences, as hypothesized. Also, the addition of price standard deviation reverses the sign on volume to negative so that, controlling for volatility, higher volume is associated with lower price differences, although this effect remains statistically insignificant. The adjusted R-squared increases from 0.17 to 0.24 when volatility is added to the regression, attesting to the marginal information value of this variable.

5.2 Market segmentation

The degree to which trading on an exchange is dominated by retail clients may be an indicator of arbitrage activity and thereby affect the extent to which its prices are integrated with those of other exchanges. We define the indicator variables $BothRetail_{i,t}$ and $OneRetail_{i,t}$ as proxies for the level of trader sophistication on the exchange. Each exchange on day t is classified as “retail” if its trade size is below its 75th percentile of trade size of the past 30 days relative to the other exchanges; otherwise it is “institutional.” $BothRetail_{i,t}$ indicates that both exchanges in pair i were classified as retail exchanges on day t , and $OneRetail_{i,t}$ indicates that one of the exchanges in pair i was classified as a retail exchange on day t . The case in which both exchanges were classified as institutional is omitted as the reference group. These indicator variables are appended to the panel regression: DESCRIP STATS RE INDICATOR VARIABLES: 1) PERSISTENCE: WE WANT TO SHOW THAT THESE INDICATORS ARE FAIRLY STABLE, AS THEY ARE IN THE DATA.

(2) EXCHANGE PAIRS THAT ARE CATEGORIZED AS ONERETAIL, BOTH RETAIL.

$$PriceDif_{i,t} = \beta_0 + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_{i,t} + \beta_5 AvgPriceSD_{i,t} + \beta_6 BothRetail_{i,t} \beta_7 OneRetail_{i,t} + \epsilon_{i,t} \quad (3)$$

Column (3) of Table 2 shows coefficients of both indicator variables are positive, indicating that relative to the relationship between two institutional pairs, there are on average larger price differences when one of the exchanges in the pair is a retail exchange, as compared to when both exchanges are institutional. The effect is stronger for *both_retail* than for *one_retail*, consistent with intuition.

5.3 Risk Premia

In addition to search frictions that result in short-run price deviations, risk premia could result in longer-lasting wedges between bitcoin prices. We had previously discussed how bitcoin exchanges are heterogenous with respect to their institutional risk such as compliance failure, risk of bankruptcy, etc. (see section 4.2). We add panel and period fixed effects to the panel regression as the fixed effects estimates incorporate these exchange-specific risks, in addition to other characteristics that are peculiar to the exchange pair:

$$PriceDif_{i,t} = \alpha_0 + \alpha_i + \alpha_t + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_{i,t} + \beta_5 AvgPriceSD_{i,t} + \beta_6 BothRetail_{i,t} \beta_7 OneRetail_{i,t} + \epsilon_{i,t} \quad (4)$$

In equation (4), α_i is the panel fixed effect and α_t is the period fixed effect. Column (4) of Table 2 shows the results from estimating the regression with exchange-pair fixed effects. To account for deterministic changes over time, column (5) of Table 2 shows results after including both exchange-pair-fixed effects and period fixed effects. The results are similar for the two cases, except that the coefficients on *both_retail* and *one_retail* become less significant when period fixed effects are added, likely reflecting the persistence of the retail/institutional indicators. The estimates of the regressors with panel fixed effects are similar to the results without fixed effects. When time effects are controlled for, the coefficients on the fee and the standard deviation are higher but with the same sign. IS THIS EFFECT STATISTICALLY SIGNIFICANT?

Table 3 shows the panel fixed effect coefficients α_i for each exchange pair, with the first column showing results without the period fixed effect and the second column including both types of fixed effects. COULD WE SHOW THIS AS FOLLOWS? THE ROWS ARE BITSTAMP COIN-

BASE ITBIT KRAKEN; COLUMNS ARE BITFINEX BTC-E AND DIFFERENCE BETWEEN THE TWO. OMITTED GROUP WOULD BE KRAKEN-COINBASE SO WE COULD HAVE ESTIMATE OF BITFINEX BITSTAMP. PANEL A: PERIOD FE; PANEL B: PERIOD+TIME. The indicator for the Bitfinex–Bitstamp pair is omitted as the reference group. WE SHOULD NOTE "OMITTED GROUP" INSTEAD OF "0" IN THE TABLE. Consistent with the descriptive statistics, the fixed effects for the exchange pairs involving BTC-e are larger in magnitude as compared to the Bitfinex-related exchange pairs. For example, in column (1) of Table 3, the estimate of the fixed effect for the Bitfinex-Coinbase pair is 0.23 while that for the BTC-e-Coinbase pair is 1.15. In other words, after controlling for transactions costs and volatility, the absolute price difference between Coinbase and BTC-e is higher by an average of 0.92 percentage points relative to the Bitfinex-Coinbase price difference. Similarly, the price difference between BTC-e-Itbit is higher than that between Bitfinex-Itbit by 0.50 percentage points, and it higher by 0.49 percentage points between BTC-e-Kraken as compared to Bitfinex-Kraken. To the extent that the fixed effects estimates are mainly picking up institutional risk factors, the results indicate that institutional factors might explain the longer-run persistent component of the price differences between BTC-e and the other exchanges. COULD WE PROVIDE RELATIVE MAGNITUDES FOR RISK PREMIA FROM ABOVE AND FRICTIONS FROM REGRESSIONS? LATTER WOULD BE MEAN OF $X \times \text{ESTIMATED COEFFICIENT ON } X$, WHERE $X = \text{SEARCH FRICTION MEASURE}$. WE SHOULD ALSO SHOW A GRAPH OF THE TIME FIXED EFFECTS IF THERE ARE INTERESTING PATTERNS.

Given the distinction between the BTC-e and Bitfinex pairs, we further estimate regressions separately for each exchange pair, grouped into two sets: exchanges paired to BTC-e, and exchanges paired to Bitfinex. ANOTHER GROUPING COULD BE BOTHRETAIL VERSUS ONERETAIL. Considering results in Tables 4, for exchanges paired to BTC-e, we find that the bid-ask spread and the price volatility are positively and significantly related to the price difference for all exchange pairs, as before. In addition, the network fee is positively related to the price difference. The coefficients on both market depth and volume are insignificant. Thus, trading frictions impede arbitrage between these exchange pairs. After accounting for arbitrage frictions, the constant term remains positive (between 4.5 and 8.5 basis points) and significant for the exchanges paired with BTC-e, implying the existence of additional sources of price differences between BTC-e related exchanges that are not accounted for by search frictions.

Turning to Table 5, which shows results for exchanges paired with Bitfinex, the effect of search frictions on the price difference is generally weaker. While the effect of the average standard deviation is positive and significant, the effects of the remaining search frictions are of inconsistent

sign and only intermittently significant. Notably, the constant is only significant for two of the four exchange pairs: Bitfinex-Kraken and Bitfinex-Bitstamp. In both cases, the magnitude is less than 4 basis points.

The regression results are in concordance with those of the descriptive statistics (section 4.1) and the qualitative discussion (5.3), in that risk premia and search frictions are both important in driving the price differences. Moreover, there is heterogeneity between the exchange pairs, with larger price differences between BTC-e related pairs being driven by both frictions and the risk of poor exchange governance and exchange failure.

6 Dynamics of Price Deviations

In this section, we turn to examining the dynamics of the price differences and how they vary across exchanges. Our descriptive statistics indicated more persistent price differences for BTC-e related exchange pairs as compared to the Bitfinex-related pairs (see sections 4.1 and 5.3). In section 6.1, we model the dynamic interactions of the price differences, volume, volatility and liquidity across exchange pairs and show that search frictions have more persistent effects on the price difference for the BTC-e related exchange pairs. More generally, the results show that the frictions have a bigger total effect (persistence times magnitude) on the price difference for the BTC-e related exchange pairs. In section 6.2, we examine the implications of more persistent price differences for dynamic market efficiency. After showing that bitcoin prices of exchange pairs are linked by a long-run equilibrium relation, we estimate how long it takes for bitcoin prices to revert to this long-run equilibrium after short-run deviations. We also compare the amount of price discovery across exchange pairs.

6.1 Dynamic interactions between price differences, illiquidity, and volatility

To formally estimate the dynamic interactions between absolute price differences, market liquidity and volatility, we estimate impulse responses of price differences to volatility, volume and liquidity. The impulse responses trace the response (in SD units) of one variable to a one-time unit SD shock to the other variable over a 10-month period. The innovations are obtained from estimating a Vector Autoregression (VAR) using endogenous variables in the following order: the network fee *AvgNetworkFee*, the intraday price standard deviation *AvgPriceSD*, log of the dollar volume *AvgUSDVol*, log of the order book depth *AvgOB*, log of the proportional bid-ask spread *AvgBASpread* and the absolute price difference *PriceDif*. The ordering builds on prior results (see

section 5.1) that transactions costs and volatility explain price differences—in other words, they are relatively more exogenous to price differences and hence come earlier in the ordering. As impulse responses are sensitive to the ordering of variables in the VAR, we estimate generalized impulse responses (Pesaran and Shin (1998)) which are not sensitive to the ordering.¹⁶ Since we reject the null hypothesis of unit roots for the variables, we estimate the VAR in levels. The number of lags is selected on the basis of information criteria.

Figure 9 shows the accumulated impulse responses of the absolute price difference to volatility and volume. Each panel has two columns, with the first column showing the responses for BTC-e paired with a particular exchange, and the second column showing responses for Bitfinex paired with the same exchange. Thus, Panel A shows responses for the BTC-e Bitstamp (column 1) and Bitfinex Bitstamp (column 2) pairs. For all exchange pairs, a one-time one SD shock to volatility or a one-time SD shock to volume results in higher absolute price differences. The accumulated response is greater for exchange pairs involving BTC-e than for pairs involving Bitfinex, mostly due to the greater persistence of the responses. For example, in Panel A, the response of the price difference to shocks in volatility is increasing even after 10 months for the BTC-e Bitstamp pair, while the response is flat after 3 months for the Bitfinex Bitstamp pair. The total response after 10 months is also higher for the BTC-e Bitstamp pair (0.73 SD versus 0.44 SD for the Bitfinex Bitstamp pair). One exception is the exchange pair involving Coinbase, in which case the Bitfinex-Coinbase response has greater persistence and a larger accumulated response (Panel B).

Figure 10 shows the impulse responses of absolute price difference to the order book depth and the bid-ask spread. A shock to the order book depth generally results in lower price differences. Similar to its responses to volume and volatility, the accumulated response of price differences to order book depth is generally greater for BTC-e related exchange pairs than Bitfinex related pairs, again mostly due to the greater persistence of responses for the BTC-e related pairs. For example, in Panel A, the response is more persistent and the accumulated response is greater for the BTC-e Bitstamp pair than for the Bitfinex Bitstamp pair. Once more, the pair involving Coinbase (Panel B) is an exception, as the response for the BTC-e Coinbase pair is not significantly different from zero whereas the response is negative and significant for the Bitfinex Coinbase pair. A shock to the bid-ask spread typically results in higher price differences. However, the accumulated responses do not follow a consistent pattern as in some cases (for example, in Panel C), they are greater for BTC-e related pairs while in other cases (for example, in Panels A and B), they are greater for Bitfinex related pairs and, in the case of Panel D, they are about the same.

¹⁶The generalized impulse responses from an innovation to the j -th variable are derived by applying a variable specific Cholesky factor computed with the j -th variable at the top of the Cholesky ordering. Since the Cholesky factor is used, these innovations are also orthogonal.

Overall, the results are consistent with higher volatility, volume and the bid-ask spread and lower order book depth leading to persistently larger absolute price differences between exchange pairs, but with heterogeneity in the effects. The dynamic relationships are generally closely related to the magnitude of search frictions between these exchanges. The responses of price differences to shocks in volatility, volume and liquidity are more persistent for BTC-e related exchange pairs (which exhibit greater search frictions) while Bitfinex related exchange pairs experience more attenuated effects on price differences. Moreover, the total effect of frictions on price differences (persistence times magnitude) is also generally greater for BTC-e related pairs.

6.2 Persistence of Price Deviations and Information Share, by Exchange Pair

Given that the asset traded in different exchanges has identical cash flows, we expect that bitcoin prices across exchange pairs should be related in a long-run equilibrium. However, frictions could result in short-run deviations of bitcoin prices from the equilibrium. Then the efficacy of arbitrage on an exchange is indicated by how long it takes for bitcoin prices to revert to the long-run equilibrium after a deviation. Further, since quicker arbitrage dynamics is likely to result in greater price discovery, we estimate the dynamics of price discovery across exchange pairs, based on the share of the variance of price changes contributed by one exchange of an exchange pair.

The speed of adjustment is estimated using a Vector Error Correction Model (VECM) using the log of prices as the endogenous variables, and lagged transactions costs, volume and volatility measures as exogenous variables, as follows:

$$\begin{aligned}
\Delta \text{Log} P_{1,t} &= \alpha_{01} + \alpha_{11}(\text{Log} P_{1,t-1} - c - d \times \text{Log} P_{2,t-1}) \\
&\quad + \alpha_{21} X_{t-1} + \sum_{k=1}^l \beta_{k1} \Delta \text{Log} P_{1,t-k} + \sum_{k=1}^l \gamma_{k1} \Delta \text{Log} P_{2,t-k} + \varepsilon_{1,t} \\
\Delta \text{Log} P_{2,t} &= \alpha_{02} + \alpha_{12}(\text{Log} P_{1,t-1} - c - d \times \text{Log} P_{2,t-1}) \\
&\quad + \alpha_{22} X_{t-1} + \sum_{k=1}^l \beta_{k2} \Delta \text{Log} P_{1,t-k} + \sum_{k=1}^l \gamma_{k2} \Delta \text{Log} P_{2,t-k} + \varepsilon_{2,t}
\end{aligned} \tag{5}$$

where P_1 is the price of bitcoin on exchange 1 (one of Bitstamp, Coinbase, Itbit or Kraken exchanges) and P_2 is the bitcoin price on exchange 2 (either BTC-e or Bitfinex). X is a vector of transactions costs, volume and volatility variables that are treated as exogenous: *AvgBAspread*, *AvgOB*, *AvgUSDVol*, *AvgPriceSD* and *AvgNetworkfee*. The number of lags l is determined using lag exclusion Wald tests. The speed of adjustment is measured as twice the half-life of a shock (where the half-life is the number of days it takes to make up 50% of the deviation from the long-

run equilibrium). We measure the speed of adjustment of bitcoin prices on exchange 1 in response to a deviation of bitcoin prices on exchange 2 as $-2 * \log(2) / \alpha_{11}$. *NS* indicates that the estimate of α_{11} is not significant. Finally, we estimate the information share of exchange 2 by variance decomposition of the log of P_1 , and report the fraction of the forecast error variance of $\log(P_1)$ that is explained by a shock to $\log(P_2)$, one day and 100 days after the shock.

Table 6 shows the results. In Panel A, P_2 is the bitcoin price on the BTC-e exchange. The first row of the table shows that all exchange prices are cointegrated of order 1 at the 0.05 percent level or lower, according to the Johansson Cointegration tests, indicating that prices of all exchange pairs are linked by a long-run equilibrium relation. The next rows show that the speed with which bitcoin prices on exchange 1 adjusts when the bitcoin price on BTC-e deviates from the long-run equilibrium relation is relatively slow, equal to 7 days on ItBit and 4 days on Kraken. α_{11} is not significantly different from zero on Bitstamp and Coinbase, indicating that bitcoin prices on these exchanges are unresponsive to movements of the bitcoin price on BTC-e. Finally, α_{12} is always insignificant, indicating the bitcoin prices on BTC-e behave as if exogenous to changes in bitcoin prices on exchange 1. The last two rows of the table shows that the information share of BTC-e ranges from being insignificantly different from 0% relative to Bitstamp to 18% relative to Kraken 100 days after the shock.

Panel B of Table 6 shows results when exchange 2 is Bitfinex. Once again, bitcoin prices are related in a long-run equilibrium for all exchange pairs. As compared to their responses to BTC-e prices, exchange 1 prices respond relatively quickly to deviations of Bitfinex prices, reverting to the long-run equilibrium in less than 2.5 days in all cases. Even bitcoin prices on Bitstamp and Coinbase respond quickly to changes in the Bitfinex price, whereas these prices were unresponsive to BTC-e price movements. Bitfinex prices, however, do not adjust to movements in bitcoin prices on exchange 1, as shown by the insignificant coefficients of α_{12} . Relative to BTC-e, more price discovery occurs on Bitfinex as its information share is 5% or more in all cases except one. For every exchange pair, the information share on Bitfinex is greater after 100 days and also for the day after the shock, except with respect to Coinbase.

In summary, we find that search frictions are associated with more persistent price differences and slower price discovery. BTC-e has larger, more persistent price differences and higher frictions relative to Bitfinex (sections 4.1 and 5.3). Consistently, bitcoin prices on exchanges adjust more slowly to deviations of BTC-e prices from the long-run equilibrium relation as compared to equilibrium deviations of Bitfinex prices. Further, there is less price discovery on BTC-e as compared to Bitfinex immediately after a shock and in the “steady state” (100 days after the shock).

7 Conclusion

In this paper, we study price differences of the virtual currency bitcoin trading on 6 exchanges constituting 15 exchange pairs where the price is quoted in US dollars (USD), with the three largest exchanges being Bitfinex, Bitstamp and BTC-e. Unlike asset pairs previously studied in the literature, bitcoin in different exchanges is, by design, the same, fully fungible, asset with identical payoffs and no currency risk. As such, bitcoin constitutes an ideal asset where the law of one price should be satisfied. Even in this ideal case, however, we show the existence of persistent, statistically significant differences between bitcoin prices in multiple USD exchanges. These price difference pairs appear to constitute two groups. In one group, where the reference exchange is BTC-e, bitcoin *consistently* trades at substantially lower prices on BTC-e compared to the other exchange. For the remaining exchange pairs, price differences are typically more *episodic* and smaller in magnitude than those in the first group, and they switch signs (i.e. the same exchange pair has a series of positive price differences followed by a series of negative differences).

We find that price differences across exchanges are related to the magnitude of arbitrage frictions. Some of these differences are related to the implicit costs of trading, such as illiquidity and return volatility. Indeed, we find that the absolute values of price differences are positively related to the bid-ask spread, order book depth and volatility and negatively related to volume. We also account for explicit fees, in particular bitcoin network transfer fees. After accounting for the implicit and explicit trading costs, we find that the average absolute price difference remains statistically significant for the BTC-e referenced exchange pairs but not for most of the remaining exchange pairs. Moreover, impulse responses indicate that shocks to illiquidity and volatility have persistent effects on price differences between BTC-e related exchange pairs, lasting 10 days or more. In contrast, the effect of shocks on price differences are more attenuated, lasting 2 to 3 days, between Bitfinex-related exchange pairs.

We further examined market segmentation and institutional factors that might determine bitcoin arbitrage profits. Segmentation may arise from a concentration of sophisticated traders on certain exchanges. We find that absolute price differences are positively related to the incidence of retail trading. We accounted for institutional factors through exchange-pair fixed effects. These fixed effects are generally statistically significant, implying that institutional features matter. Other factors, such as anonymity and exchange failure risk are also important, although we cannot quantify them. In particular, BTC-e offers enhanced anonymity relative to other exchanges, making it potentially attractive to users wishing to convert illicitly-obtained bitcoin into fiat currency, thus creating increased selling pressure (and hence lower prices) compared to other exchanges.

The results of Vector Error Correction Model (VECM) suggest that the prices of all exchange pairs are cointegrated of order one, indicating the existence of a long-run equilibrium relation between prices of each exchange pair. However, there are substantial differences between exchanges in the speed of adjustment of prices to deviations from the long-run equilibrium relation. Specifically, the half-life of a price deviation (i.e. the number of days required to eliminate half of the price deviation) is two to three times higher for BTC-e, as compared to other exchanges. Moreover, the information share of BTC-e is small compared to other exchanges, indicating that little price discovery occurs on this exchange. By contrast, a substantial part of price discovery occurs on Bitfinex, relative to other exchanges. These results imply that the speed of arbitrage is related to the arbitrage frictions.

We contribute to the literature by showing that microstructure frictions, market segmentation and institutional features all contribute to limiting arbitrage between bitcoin exchanges. Thus, limits to arbitrage remain relevant even in the context of a homogeneous asset class where many of the frictions in more traditional asset markets are absent.

Table 1: Exchange-Related Fees

The table shows typical fees as of August 2016 for the major bitcoin-USD exchanges. Deposit and withdrawal fees are for wire transfers to deposit or withdraw USD from the exchange. Trading fees are different from “taker” (the side that crosses the bid-ask spread) and “maker” (the side whose order is filled). The data is obtained from the exchange websites.

Exchange	USD Deposit Fee	USD Withdrawal Fee	“Taker” Transaction Fees	“Maker” Transaction Fees
Bitfinex	0.10%	0.10%	0.10-0.20%	0.00-0.10%
Bitstamp	0.05%	0.09%	0.10-0.25%	0.10-0.25%
Coinbase	0	0	0.10-0.25%	0.00%
BTC-e	\$20		0.20-0.50%	0.20-0.50%
Kraken	0.19%	0.19%	0.10-0.26%	0-0.16%
itBit	\$10	\$20	0.20%	0.00%

Table 2: Panel Regression Results on Determinants of Bitcoin Price Deviations

The table shows the results from the following regression:

$$PriceDif_{i,t} = \alpha_0 + \alpha_i + \alpha_t + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_t + \beta_5 AvgPriceSD_{i,t} + \beta_6 BothRetail_{i,t} + \beta_7 OneRetail_{i,t} + \varepsilon_{i,t}$$

where $PriceDif_{i,t}$ for an exchange pair i on day t is the absolute value of the bitcoin price difference between the exchanges, divided by the the average of the bitcoin prices on the two exchanges. The exchanges are Bitfinex, Bitstamp, BTC-e, Coinbase, Kraken and ItBit. $LnAvgBASpread_{i,t}^1$ is the the bid-ask spread divided by the mid-quote. $LnAvgOB_{i,t}^1$ is the daily sum of orders in the order book within 1% of the average daily price, divided by the daily volume. $LnAvgUSDVol_{i,t}$ is the trading volume in millions of USD. $LnAvgNetworkFee_t$ is the average transaction fee (in USD) for all transactions added to the blockchain. $AvgPriceSD_{i,t}$ is the intraday price standard deviation normalized by the daily price. $BothRetail_{i,t}$ indicates that both exchanges in pair i were classified as retail exchanges on day t , and $OneRetail_{i,t}$ indicates that one of the exchanges in pair i was classified as a retail exchange on day t . α_i is the panel fixed effect. α_t is the period fixed effect. $\varepsilon_{i,t}$ is an error term.

	(1)	(2)	(3)	(4)	(5)
Log(Avg. spread / price)	0.388*** (0.0464)	0.222*** (0.0439)	0.213*** (0.0376)	0.372*** (0.0389)	0.321*** (0.0436)
Log(Avg. orders $\pm 1\%$ of price)	-0.223*** (0.0503)	-0.176** (0.0524)	-0.116** (0.0330)	-0.0295 (0.0278)	-0.0224 (0.0187)
Log(Avg. USD volume (Mil. USD))	0.0224 (0.0757)	-0.176 (0.0895)	-0.0720 (0.0647)	-0.120* (0.0535)	-0.121 (0.0586)
Log(Avg. Network Transaction Fee)	0.226 (0.142)	0.268 (0.151)	0.325* (0.141)	0.415** (0.102)	1.629** (0.537)
Avg. intraday price std. dev.		39.59*** (6.312)	38.37*** (5.893)	42.25*** (5.113)	90.78*** (8.263)
Both retail			0.793*** (0.182)	-0.0136 (0.107)	0.104 (0.156)
One retail			0.409** (0.126)	0.0977* (0.0396)	0.154* (0.0621)
Constant	2.426* (0.824)	1.668 (0.786)	1.966* (0.687)	4.162*** (0.642)	7.604*** (2.096)
Observations	9,396	9,380	9,380	9,380	9,380
R-Squared	0.165	0.243	0.286	0.366	0.555
Std. Errors	Clustered By Exch. Pair	Clustered By Exch. Pair	Clustered By Exch. Pair	Clustered By Exch. Pair	Clustered By Exch. Pair
Fixed Effects	None	None	None	Panel	Panel, Time

Standard errors in parentheses

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

Table 3: Panel Fixed Effects Estimates from Panel Regressions

The table shows the panel fixed effects estimates from the following regression:

$$PriceDif_{i,t} = \alpha_0 + \alpha_i + \alpha_t + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_t + \beta_5 AvgPriceSD_{i,t} + \beta_6 BothRetail_{i,t} \beta_7 OneRetail_{i,t} + \epsilon_{i,t}$$

where $PriceDif_{it}$ for an exchange pair i on day t is the absolute value of the bitcoin price difference between the exchanges, divided by the the average of the bitcoin prices on the two exchanges. α_i is the panel fixed effect. α_t is the period fixed effect. The exchanges are Bitfinex, Bitstamp, BTC-e, Coinbase, Kraken and ItBit. The regressors are defined in Table 2.

	(1)	(2)
Bitfinex, Bitstamp	0	0
	(.)	(.)
Bitfinex, Coinbase	0.228**	0.133
	(0.0746)	(0.0897)
Bitfinex, Itbit	0.213***	0.252***
	(0.0204)	(0.0184)
Bitfinex, Kraken	-0.193*	-0.180
	(0.0728)	(0.0877)
Bitstamp, Coinbase	-0.0411	-0.0869
	(0.0570)	(0.0915)
Bitstamp, Itbit	-0.107*	-0.0408
	(0.0363)	(0.0377)
Bitstamp, Kraken	-0.418***	-0.396***
	(0.0755)	(0.0772)
Btce, Bitfinex	0.920***	0.872***
	(0.0413)	(0.0604)
Btce, Bitstamp	0.530***	0.509***
	(0.0381)	(0.0674)
Btce, Coinbase	1.148***	1.046***
	(0.111)	(0.178)
Btce, Itbit	0.711***	0.727***
	(0.0585)	(0.0945)
Btce, Kraken	0.296***	0.270
	(0.0574)	(0.128)
Itbit, Coinbase	0.0322	0.0142
	(0.0722)	(0.114)
Kraken, Coinbase	-0.431***	-0.445**
	(0.0530)	(0.130)
Kraken, Itbit	-0.420**	-0.350**
	(0.106)	(0.0986)
Constant	4.162***	7.604**
	(0.642)	(2.096)
Observations	9,380	9,380
R-Squared	0.366	0.555
Std. Errors	Clustered	Clustered
	By Exch. Pair	By Exch. Pair
Time Fixed Effects?	No	Yes

Standard errors in parentheses

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

Table 4: Determinants of Bitcoin Price Deviations: Exchanges Paired with BTC-e

The table shows results from the following regression:

$$PriceDif_{i,t} = \beta_0 + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_t + \beta_5 AvgPriceSD_{i,t} + \varepsilon_{i,t}$$

where $PriceDif_{it}$ for an exchange pair i on day t is the absolute value of the bitcoin price on exchange X minus its price on BTC-e, divided by the the average of the bitcoin prices on the two exchanges. X =Bitfinex, Bitstamp, Coinbase, Kraken, ItBit. All regressors, except the standard deviation $AvgPriceSD$, are in log form. Further, regressors are calculated for each exchange and then averaged over the two exchanges in pair i . $LnAvgBASpread_{it}^1$ is the the bid-ask spread divided by the mid-quote. $LnAvgOB_{it}^1$ is the daily sum of orders in the order book within 1% of the average daily price, divided by the daily volume. $LnAvgUSDVol_{it}$ is the trading volume in millions of USD. $LnAvgNetworkFee_t$ is the average transaction fee (in USD) for all transactions added to the blockchain. $AvgPriceSD_{i,t}$ is the intraday price standard deviation normalized by the daily price. $\varepsilon_{i,t}$ is an error term.

	(1)	(2)	(3)	(4)	(5)
Log(Avg. spread / price)	0.472*** (0.137)	0.372** (0.117)	0.619*** (0.117)	0.358*** (0.0824)	0.582*** (0.160)
Log(Avg. orders $\pm 1\%$ of price)	0.202 (0.117)	-0.0169 (0.124)	-0.118 (0.128)	-0.0486 (0.0537)	-0.112 (0.0905)
Log(Avg. USD volume (Mil. USD))	0.272* (0.138)	0.178 (0.188)	0.196 (0.198)	-0.407* (0.184)	0.127 (0.161)
Avg. intraday price std. dev.	47.15*** (12.59)	24.51* (11.58)	24.50 (13.13)	67.36*** (15.81)	34.03** (11.73)
Log(Avg. Network Transaction Fee)	0.673** (0.225)	0.325 (0.217)	0.704* (0.290)	1.304*** (0.327)	0.619* (0.269)
Constant	8.423*** (1.825)	4.423** (1.488)	7.322*** (1.816)	7.590*** (1.636)	6.448*** (1.820)
Observations	706	695	549	659	663
1st Exchange	BTC-e	BTC-e	BTC-e	BTC-e	BTC-e
2nd Exchange	Bitfinex	Bitstamp	Coinbase	Kraken	Itbit
Std. Errors	Newey-West Max lag=3	Newey-West Max lag=3	Newey-West Max lag=3	Newey-West Max lag=3	Newey-West Max lag=3

Standard errors in parentheses

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

Table 5: Determinants of Bitcoin Price Deviations: Exchanges Paired with Bitfinex

The table shows results from the following regression:

$$PriceDif_{i,t} = \beta_0 + \beta_1 LnAvgBASpread_{i,t} + \beta_2 LnAvgOB_{i,t} + \beta_3 LnAvgUSDVol_{i,t} + \beta_4 LnAvgNetworkFee_t + \beta_5 AvgPriceSD_{i,t} + \epsilon_{i,t}$$

where $PriceDif_{it}$ for an exchange pair i on day t is the absolute value of the bitcoin price on exchange X minus its price on Bitfinex, divided by the the average of the bitcoin prices on the two exchanges. X =Bitfinex, Bitstamp, Coinbase, Kraken, ItBit. All regressors, except the standard deviation $AvgPriceSD$, are in log form. Further, regressors are calculated for each exchange and then averaged over the two exchanges in pair i . $LnAvgBASpread_{it}^1$ is the the bid-ask spread divided by the mid-quote. $LnAvgOB_{it}^1$ is the daily sum of orders in the order book within 1% of the average daily price, divided by the daily volume. $LnAvgUSDVol_{it}$ is the trading volume in millions of USD. $LnAvgNetworkFee_t$ is the average transaction fee (in USD) for all transactions added to the blockchain. $\epsilon_{i,t}$ is an error term.

	(1)	(2)	(3)	(4)
Log(Avg. spread / price)	0.147** (0.0553)	0.111 (0.101)	0.318*** (0.0589)	-0.145 (0.124)
Log(Avg. orders $\pm 1\%$ of price)	0.215* (0.0848)	0.298 (0.162)	-0.0141 (0.0415)	-0.264*** (0.0797)
Log(Avg. USD volume (Mil. USD))	0.0250 (0.0553)	0.160* (0.0808)	-0.398** (0.147)	-0.216** (0.0819)
Avg. intraday price std. dev.	41.00** (12.48)	28.91* (12.27)	83.44*** (18.79)	42.31*** (11.17)
Log(Avg. Network Transaction Fee)	0.0194 (0.157)	0.271 (0.139)	0.405 (0.210)	0.0739 (0.207)
Constant	3.047* (1.292)	4.534 (2.441)	3.685*** (1.106)	-2.481 (1.500)
Observations	701	555	666	670
1st Exchange	Bitfinex	Bitfinex	Bitfinex	Bitfinex
2nd Exchange	Bitstamp	Coinbase	Kraken	Itbit
Std. Errors	Newey-West Max lag=3	Newey-West Max lag=3	Newey-West Max lag=3	Newey-West Max lag=3

Standard errors in parentheses

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

Table 6: Speed of adjustment to equilibrium deviations and price discovery

The table reports the speed of adjustment to deviations in bitcoin prices between exchange pairs and price discovery on each of the exchanges, estimated from a Vector Error Correction Model (VECM):

$$\begin{aligned}\Delta \text{Log} P_{1,t} &= \alpha_{01} + \alpha_{11}(\text{Log} P_{1,t-1} - c - d \times \text{Log} P_{2,t-1}) \\ &\quad + \alpha_{21} X_{t-1} + \sum_{k=1}^l \beta_{k1} \Delta \text{Log} P_{1,t-k} + \sum_{k=1}^l \gamma_{k1} \Delta \text{Log} P_{2,t-k} + \varepsilon_{1,t} \\ \Delta \text{Log} P_{2,t} &= \alpha_{02} + \alpha_{12}(\text{Log} P_{1,t-1} - c - d \times \text{Log} P_{2,t-1}) \\ &\quad + \alpha_{22} X_{t-1} + \sum_{k=1}^l \beta_{k2} \Delta \text{Log} P_{1,t-k} + \sum_{k=1}^l \gamma_{k2} \Delta \text{Log} P_{2,t-k} + \varepsilon_{2,t}\end{aligned}$$

Exchange 1 is one of Bitstamp, Coinbase, ItBit or Kraken exchanges. X is a vector of exogenous variables (the standard deviation of returns, the network trading fee and logs of the bid-ask spread, the order book depth and the dollar volume). l is the number of lags. In Panel A, P_1 is the bitcoin price on exchange 1 and P_2 is the bitcoin price on the BTC-e exchange. In Panel B, P_1 is the bitcoin price on exchange 1 and P_2 is the bitcoin price on the Bitfinex exchange. *Cointegrated of Order 1* is YES if there is exactly one cointegrating equation at the 0.05 percent level, according to the Jonansen Cointegration tests. *Speed of adjustment* of exchange 1 to BTC-e or to Bitfinex is $-2 * \log(2) / \alpha_{11}$. *NS* indicates that the estimate of α_{11} or α_{12} is not significant. The *information share* of BTC-e or Bitfinex is its share in the variance decomposition of the exchange 1 bitcoin price using the Cholesky decomposition. Estimates are shown for one day and 100 days after the shock. *S.E.* indicates the standard error.

(a) Panel A: BTC-e relative to other exchanges

	Exchange 1			
	Bitstamp	Coinbase	ItBit	Kraken
Cointegrated of Order 1?	YES	YES	YES	YES
Cointegration coefficient α_{11}				
Estimate	-0.03	-0.07	-0.21	-0.38
T-statistics	-0.30	-0.69	-2.10	-4.04
Speed of adjustment, exchange 1 to BTC-e	NS	NS	6.59	3.65
Cointegration coefficient α_{12}				
Estimate	0.14	0.06	0.08	0.08
T-stat	1.40	0.61	0.80	0.89
Information share (%) of BTC-e				
After day 1	0.00	1.49	1.98	5.26
S.E.	0.04	0.04	0.04	0.03
After day 100	0.18	2.22	7.00	18.04
S.E.	0.29	0.22	0.29	0.30

(b) Panel B: Bitfinex relative to other exchanges

	Exchange 1			
	Bitstamp	Coinbase	ItBit	Kraken
Cointegrated of Order 1?	YES	YES	YES	YES
Cointegration coefficient α_{11}				
Estimate	-0.93	-0.61	-0.90	-0.56
T-statistics	-3.45	-1.99	-4.43	-4.76
Speed of adjustment, exchange 1 to Bitfinex	1.49	2.29	1.54	2.46
Cointegration coefficient α_{12}				
Estimate	-0.33	0.01	-0.09	-0.03
T-stat	-1.18	0.03	-0.41	-0.21
Information share (%) of Bitfinex				
After day 1	6.35	0.88	10.48	12.17
S.E.	0.04	0.04	0.04	0.04
After day 100	13.23	5.05	14.45	28.12
S.E.	0.33	0.27	0.32	0.33

Figure 1: Average Price of Bitcoin in USD: 2012-2016

The figure shows the time series of the bitcoin index price. The data source is bitcoinaverage.com and is downloaded from [Quandl](https://www.quandl.com/). The average price is calculated as the trading volume weighted average across several currency markets, converted to USD at fiat exchange rates. The sample period is January 2012 to August 2016.




Figure 2: Total volume in the bitcoin–dollar exchange market: 2012-2016

The figures show the 60-day moving average of the daily volume in bitcoin exchanges, in units of bitcoin (Panel a) and in US dollars (Panel b). The bitcoin exchanges are Bitfinex, Bitstamp BTC-e, Coinbase, itBit, MtGox, Kraken and other USD-BTC exchanges. The data source is bitcoincharts.com. The sample period is January 2012 to August 2016.

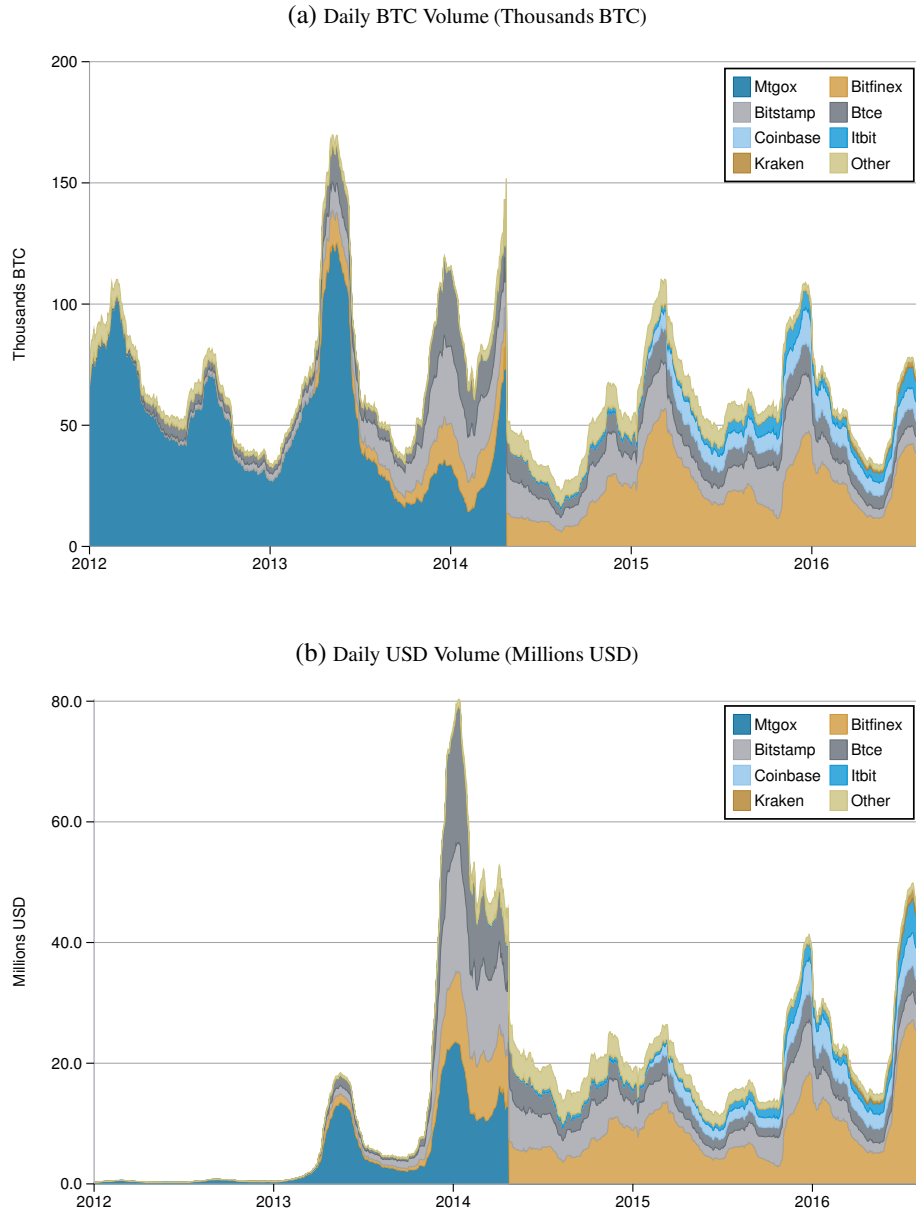


Figure 3: Deviations in Bitcoin Prices Relative to BTC-e

The figure plots the daily time series of differences between the price of bitcoin on an exchange X minus its price on BTC-e, normalized by the average of the prices on the two exchanges, where $X=(\text{Bitfinex}, \text{Bitstamp}, \text{Coinbase}, \text{itBit})$. The price is the volume-weighted average price of the day. The price difference is capped in the range -10% to $+10\%$ but the maximum absolute difference is indicated below each chart. The magnitude of the normalized price difference is shown as a blue area plot. The orange line shows the average signed price differences for the respective exchange pair. The dark red line shows the average absolute price difference between the exchanges. The data source is bitcoincharts.com. The horizontal time axis varies based on data availability.

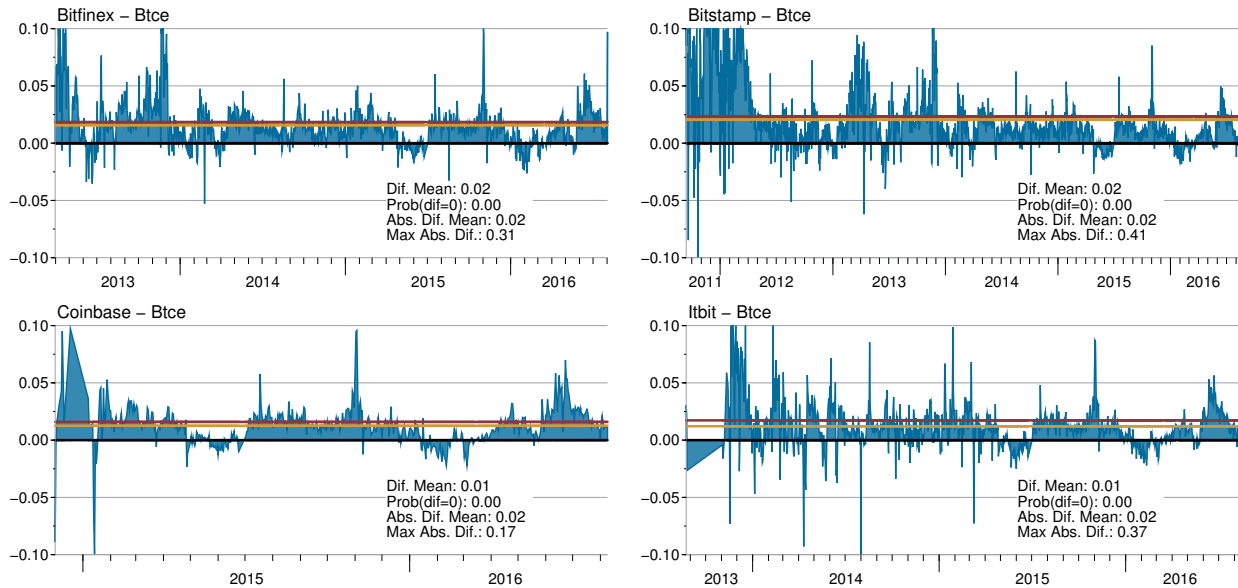


Figure 4: Deviations in Bitcoin Prices Relative to Bitfinex

The figure plots the daily time series of differences between the price of bitcoin on an exchange X minus its price on Bitfinex, normalized by the average of the prices on the two exchanges, where $X=(\text{Bitstamp}, \text{BTC-e}, \text{Coinbase}, \text{itBit})$. The price is the volume-weighted average price of the day. The price difference is capped in the range -10% to $+10\%$ but the maximum absolute difference is indicated below each chart. The magnitude of the normalized price difference is shown as a blue area plot. The orange line shows the average signed price differences for the respective exchange pair. The dark red line shows the average absolute price difference between the exchanges. The data source is bitcoincharts.com. The horizontal time axis varies based on data availability.

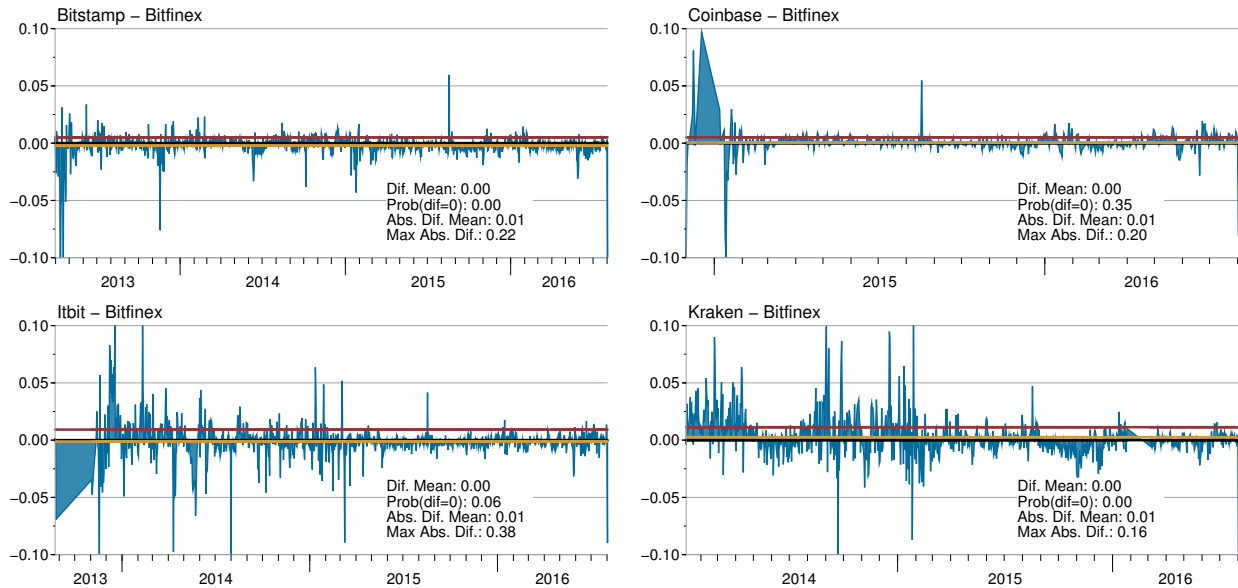


Figure 5: Mechanics of a Bitcoin Arbitrage Strategy

The diagram shows the execution of a possible arbitrage strategy involving buying bitcoin on BTC-e and selling it on Bitfinex. The data on fees are from the exchange websites, as of September 2016.

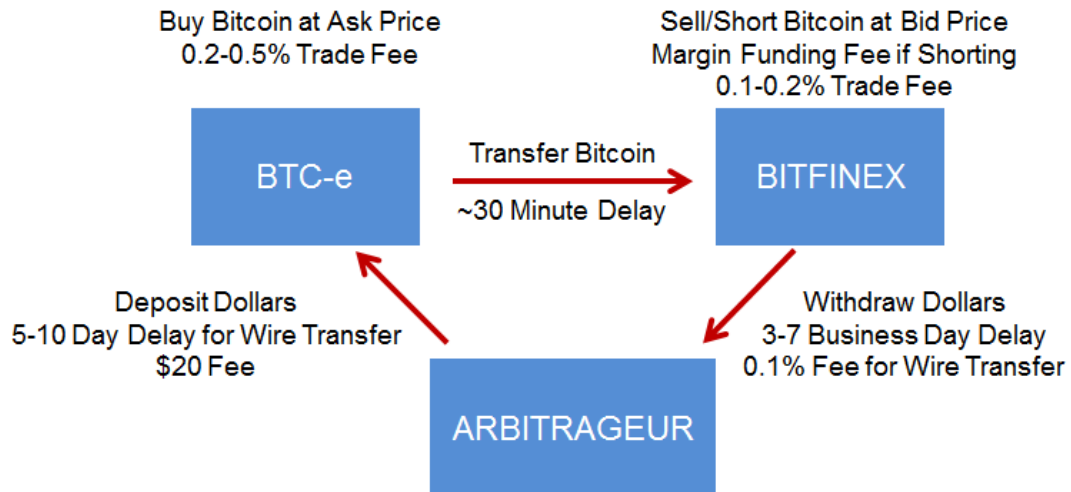


Figure 6: Average bitcoin network fee per transaction (in converted USD): 2013-2016

The figure shows the average per transaction of bitcoin network transaction fees that are used to incentivize miners to include the transaction in their blocks. The high average fee in April 2016 was likely due to an error when a user specified a transaction fee of 316 BTC, or roughly \$147,000 at the time ([Coindesk](#)). The data is obtained from [Blockchain.info](#).




Figure 7: Intraday Volatility of Bitcoin Prices

The figure shows the intraday volatility of the bitcoin price index and the absolute value of the difference in prices of bitcoin trading on Bitfinex and BTC-e. The data is obtained from [bitcoincharts.com](#).




Figure 8: Distribution of transaction sizes of USD–bitcoin trades, by exchange: Jan-Aug 2016

The figure shows box-whisker plots of the distribution of transactions sizes of trades (denominated in bitcoins) in all Bitcoin-USD exchanges in 2016. The middle line within the box represents the median value. The top and bottom of the box represent the 75th and 25th percentiles of transaction size, respectively. The top and bottom of the “whiskers” are the upper and lower adjacent values, defined as the most extreme values within 1.5 times the interquartile range (the 75th minus the 25th percentiles). The sample period is January to August of 2016. The data used is transaction-level data obtained from bitcoincharts.com.


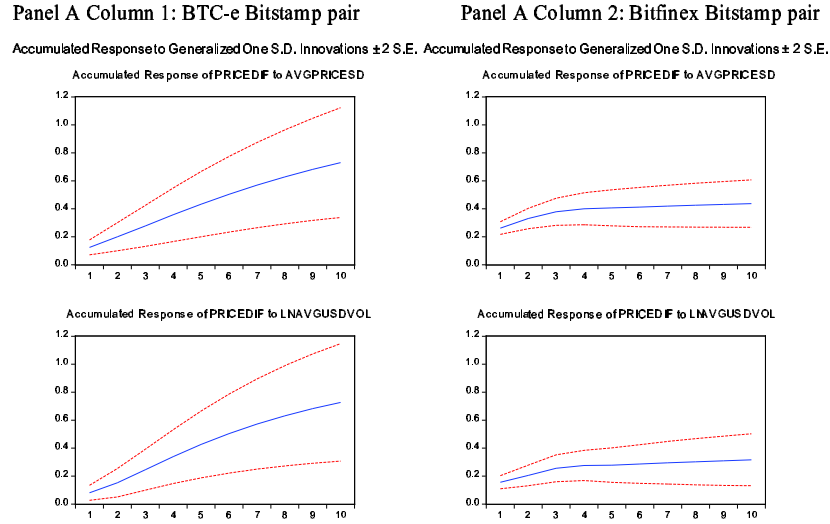


Figure 9: Impulse responses of price differences, volume, and volatility: : BTC-e versus Bitfinex

We estimate a Vector Autoregression (VAR) using endogenous variables in the following order: the network fee *AvgNetworkFee*, the intraday price standard deviation *AvgPriceSD*, log of the dollar volume *AvgUSDVol*, log of the sum of bid and ask orders normalized by the daily volume *AvgOB*, log of the proportional bid-ask spread *AvgBAspread*, the absolute price difference *PriceDif*. All variables other than the price difference is averaged over the respective exchange pairs. The figures show impulse response functions, along with ± 2 standard error bands, of *PriceDif* to the log of *AvgUSDVol* and *AvgPriceSD*. In each panel, we pair BTC-e and Bitfinex with one of the following exchanges: Bitstamp (Panel A), Coinbase (Panel B), ItBit (Panel C) and Kraken (Panel D).

(a) Panel A: Bitstamp







(b) Panel B: Coinbase

Panel B Column 1: Btce Coinbase pair

Panel B Column 2: Bitfinex Coinbase pair

Accumulated Response to Generalized One S.D. innovations ± 2 S.E. Accumulated Response to Generalized One S.D. innovations ± 2 S.E.







(c) Panel C: itBit

Panel C Column 1: BTCE-Itbit pair

Panel C Column 2: Bitfinex-Itbit pair

Accumulated Response to Generalized One S.D. innovations ± 2 S.E. Accumulated Response to Generalized One S.D. innovations ± 2 S.E.



(d) Panel D: Kraken

Panel D Column 1: BTCE-Kraken pair

Panel D Column 2: Bitfinex-Kraken pair

Accumulated Response to Generalized One S.D. Innovations ± 2 S.E. Accumulated Response to Generalized One S.D. Innovations ± 2 S.E.





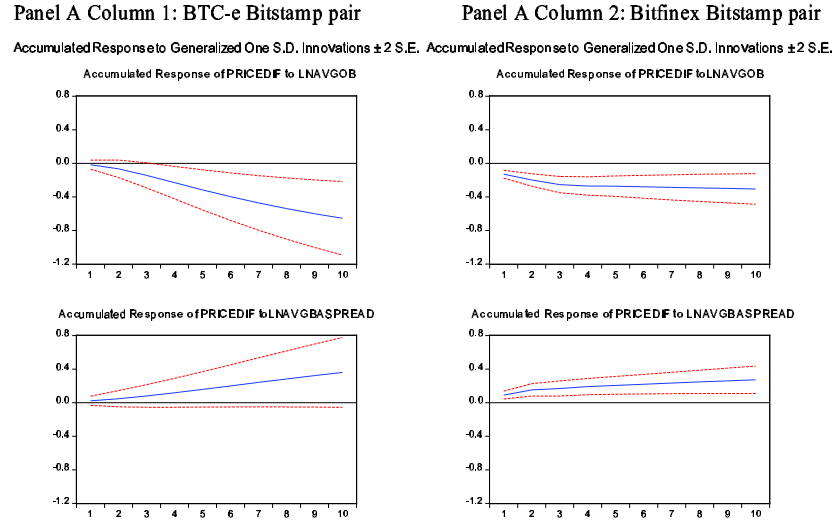


Figure 10: Impulse responses of price differences and liquidity: BTC-e versus Bitfinex

We estimate a Vector Autoregression (VAR) using endogenous variables in the following order: the network fee *AvgNetworkFee*, the intraday price standard deviation *AvgPriceSD*, log of the dollar volume *AvgUSDVol*, log of the sum of bid and ask orders normalized by the daily volume *AvgOB*, log of the proportional bid-ask spread *AvgBAspread*, the absolute price difference *PriceDif*. All variables other than the price difference is averaged over the respective exchange pairs. The figures show impulse response functions, along with ± 2 standard error bands, of *PriceDif* to the log of *AvgOB* and the log of *AvgBAspread*. In each panel, we pair BTC-e and Bitfinex with one of the following exchanges: Bitstamp (Panel A), Coinbase (Panel B), ItBit (Panel C) and Kraken (Panel D).

(a) Panel A: Bitstamp

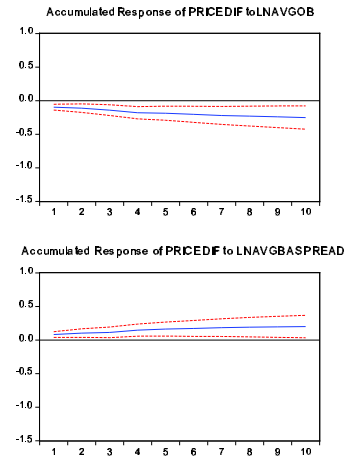
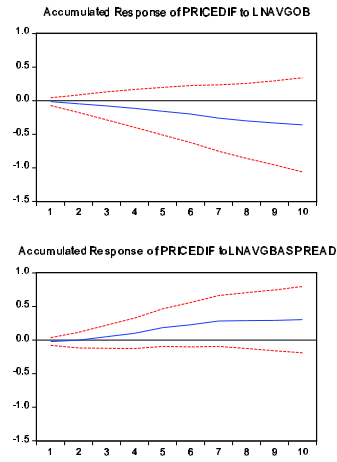


(b) Panel B: Coinbase

Panel B Column 1: Btce Coinbase pair

Panel B Column 2: Bitfinex Coinbase pair

Accumulated Response to Generalized One S.D. Innovations ± 2 S.E. Accumulated Response to Generalized One S.D. Innovations ± 2 S.E.







(c) Panel C: itBit

Panel C Column 1: BTCE-Itbit pair

Panel C Column 2: Bitfinex-Itbit pair

Accumulated Response to Generalized One S.D. Innovations ± 2 S.E. Accumulated Response to Generalized One S.D. Innovations ± 2 S.E.







(d) Panel D: Kraken

Panel D Column 1: BTCE-Kraken pair

Panel D Column 2: Bitfinex-Kraken pair


Accumulated Response to Generalized One S.D. Innovations ± 2 S.E. Accumulated Response to Generalized One S.D. Innovations ± 2 S.E.



A Additional Figures

Figure A1: History of the block size on the blockchain, 2012-2016

The figure shows a time series of the block size of successfully mined blocks, averaged at the daily level. The block size increases depending on the number of transactions included in the block.¹⁷ The block size limit is fixed at 1 megabyte—a threshold that the average block size has been steadily approaching. The data is taken from [Blockchain.info](https://blockchain.info)



B Data Appendix

Table B1: Range of Data Available on the Six Major Exchanges Considered from bitcoincharts.com

Exchange	First Date	Last Date
Bitfinex	03/31/13	08/02/16
Bitstamp	09/13/11	08/09/16
BTC-e	08/14/11	08/09/16
Coinbase	12/01/14	08/09/16
itBit	08/24/13	08/09/16
Kraken	01/07/14	08/09/16

Table B2: Range of Data Available on Other USD–BTC Exchanges from bitcoincharts.com

Exchange	First Date	Last Date	Exchange	First Date	Last Date
Anxhk	08/20/2013	10/10/2015	Crytr	05/27/2013	07/16/2014
B2C	04/01/2011	01/21/2012	Exchb	06/15/2011	10/16/2011
B7	06/14/2011	10/05/2011	Exmo	04/14/2016	07/10/2016
Bclr	12/31/2010	08/19/2011	Fbtc	06/25/2013	10/20/2013
Bcmbm	02/04/2011	08/05/2011	Global	09/26/2011	10/07/2011
Bcmlr	06/04/2010	08/03/2011	Hitbtc	12/27/2013	08/09/2016
Bcmpp	04/25/2010	06/04/2011	Ibwt	09/27/2014	11/04/2015
Bitalo	03/03/2014	08/05/2014	Icbt	03/18/2012	05/23/2014
Bitbay	05/16/2014	08/09/2016	Imcex	07/17/2011	10/14/2012
Bitbox	05/07/2013	09/24/2013	Intrsn	07/21/2011	10/20/2012
Bitcurex	04/20/2016	08/03/2016	Itbit	08/24/2013	08/09/2016
Bitfinex	03/31/2013	08/02/2016	Just	10/25/2013	10/29/2014
Bitfloor	05/23/2012	04/17/2013	Kraken	01/07/2014	08/09/2016
Bitkonan	07/02/2013	08/06/2016	Lake	03/01/2014	07/19/2015
Bitmarket	05/17/2011	12/20/2012	Localbtc	03/11/2013	08/09/2016
Bitme	07/08/2012	03/28/2013	Lybit	01/29/2013	06/08/2013
Bitstamp	09/13/2011	08/09/2016	Mtgox	07/17/2010	02/24/2014
Btc24	05/14/2012	04/13/2013	Ripple	02/10/2013	11/04/2014
Btce	08/14/2011	08/09/2016	Rock	11/12/2011	08/09/2016
Btcex	02/04/2011	07/22/2012	Ruxum	06/30/2011	09/12/2012
Btctree	05/01/2012	07/10/2012	Th	06/08/2011	02/13/2012
Cbx	07/05/2011	07/23/2016	Thlr	06/10/2011	02/13/2012
Coinbase	12/01/2014	08/09/2016	Vcx	12/25/2011	08/07/2016
Cotr	12/21/2013	01/31/2016	Weex	06/01/2012	11/26/2013
Cryptox	11/10/2011	11/19/2012	X1Coin	03/09/2014	04/04/2015

References

- Ali, Robleh, John Barrdear, Roger Clews, and James Southgate (2014) “The economics of digital currencies,” *Bank of England Quarterly Bulletin*.
- Amihud, Y. and H. Mendelson (1991) “Liquidity, Maturity, and the Yield on U.S. Treasury Securities,” *Journal of Finance*, Vol. 46, pp. 479–486.
- Caffyn, Grace (2015) “Bitcoin on the Dark Web: The Facts,” *CoinDesk*.
- Debora, E. and K. Froot (1999) “How are Stock Prices Affected by the Location of Trade?” *Journal of Financial Economics*, Vol. 53, pp. 189–216.
- Duffie, Darrell (2010) “Presidential Address: Asset Price Dynamics with Slow-Moving Capital,” *The Journal of Finance*, Vol. 65, pp. 1237–1268.
- Gagnon, Louis and George Andrew Karolyi (2010) “Multi-Market Trading and Arbitrage,” *Journal of Financial Economics*, Vol. 97, pp. 53–80.
- Gromb, Denis and Dimitri Vayanos (2010) “Limits of Arbitrage,” *Annual Reviews of Financial Economics*, Vol. 2, pp. 251–275.
- Jong, Abe De, Leonard Rosenthal, and Mathijs A. Van Dijk (2009) “The Risk and Return of Arbitrage in Dual-Listed Companies,” *Review of Finance*, Vol. 13, pp. 495–520.
- Krishnamurthy, A. (2009) “The bond/old-bond spread,” *Journal of Financial Economics*, Vol. 66, pp. 463–506.
- Lamont, Owen and Richard Thaler (2003) “Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs,” *Journal of Political Economy*, Vol. 111(2), pp. 227–268.
- McMillan, Robert (2014) “The Inside Story of Mt. Gox, Bitcoin’s \$460 Million Disaster,” *Wired*.
- Mitchell, M., T. Pulvino, and E. Stafford (2002) “Limited Arbitrage in Equity Markets,” *Journal of Finance*, Vol. 57, pp. 551–584.
- Moore, Tyler and Nicolas Christin (2013) “Beware the Middle Man: Empirical Analysis of Bitcoin-Exchange Risk,” *Financial Cryptography and Data Security*, pp. 25–33.
- Pesaran, Hashem M. and Yongcheol Shin (1998) “Impulse Response Analysis in Linear Multivariate Models,” *Economics Letters*, Vol. 75, pp. 335–346.

- Pieters, Gina and Sofia Vivanco (2016) “Bitcoin: A Hub of Criminal Activity?” *Working Paper*.
- Rosenthal, L. and C Young (1990) “The Seemingly Anomalous Price Behavior of Royal Dutch/Shell and Unilever,” *Journal of Financial Economics*, Vol. 26, pp. 123–141.
- Russell, Jon (2015) “Bitstamp Suspends Its Service After Hackers Snatch Nearly \$5M in Bitcoins,” *TechCrunch*.
- Shleifer, Andrei and Robert Vishny (1997) “The Limits of Arbitrage,” *The Journal of Finance*, Vol. 52(1), pp. 35–55.
- Tepper, Fitz (2016) “Hacked Bitcoin Exchange Bitfinex will reduce balances by 36% to distribute losses amongst all users,” *TechCrunch*.
- Warga, A. (1992) “Bond Returns, Liquidity and Missing Data,” *Journal of Financial and Quantitative Analysis*, Vol. 27, pp. 605–617.