Cryptocurrency Exchange Arbitrage: an Investigation into Trading Profits and Market Efficiency

Daniel Benson

University of Sussex Business School

Supervisor: Prof. Carol Alexander

Presented to the Department of Accounting and Finance in partial fulfilment of the requirements for the degree of Master of Science in Financial Risk and Investment Analysis

August 2019

Abstract

Until recently, the deviations in the price of Bitcoin across exchanges had become low and stable; however, in the months following June 2019, we see large recurring arbitrage opportunities. We test trading strategies on 10-minute Bitcoin data across five cryptocurrency exchanges and find a traditional pairs trading strategy exploits these price differences best. Furthermore, we test the relationship between prices deviations across exchanges and realized volatility and volume, finding volatility plans a significant role in determining the spreads.

Keywords: Bitcoin, Arbitrage, Market efficiency, Exchange Arbitrage, Trading

Acknowledgements

Firstly, I would like to say a big thank you to my supervisor Carol Alexander who supported my research all the way through. I would also like to thank Leah Coughlan for being more than just a proof reader and supporting me through my studies. Finally, I'd like to thank Hamish Massie for keeping me company through those long days in the library and for overall support.

Table of Contents

List of Tables

List of Figures

1 Introduction

Since the introduction of Bitcoin's white paper (Nakamoto 2008), the market for cryptocurrencies has evolved dramatically. Today, we see more than 100 exchanges worldwide, around 2,322 cryptocurrencies and over 230 venture capital funds working within crypto. Therefore, it is no surprise that investors and hedge funds are searching for new ways to profit in this innovative market. It does not take much investigation to see that there are significant deviations in the price of Bitcoin across exchanges. With researchers such as Makarov and Schoar (2019) suggesting there are upwards of $900 million worth of profits available as well as phenomenon such as the "Kimchi premium",[1] exchange arbitrage may be a reasonable place to look. With large price deviations between exchanges and the lack of regulation, it not only increases the role of arbitrageurs, but it also makes it an easy place to derive profit. The introduction of previously used arbitrage strategies such as pairs trading if implemented on cryptocurrencies may derive profits.

Pairs trading was first introduced in the 1980s by a Wall Street quant Nunzio Tartaglia, who brought mathematicians, physicists and computer scientists to financial markets to finds ways to exploit the markets. The pairs trading strategy was introduced as a rule-based strategy to exploit two stocks that are cointegrated, profiting when the spread between the stocks widens and thus converges. Today, replacing the stocks with Bitcoin on different exchanges may be a simple way to exploit cryptocurrency markets.

With such a growing market, it is no surprise that Bitcoin has become popular with academics and hedge funds. For several reasons, cryptocurrency markets make an ideal laboratory to study market efficiency and provide a suitable market for hedge funds to exploit. Cryptocurrency markets are constituted of many non-integrated exchanges, which function similar to an equity market where investors submit buy and sell orders and trades clear based on a central order book. The markets lack the regulation of a traditional equity market; therefore, the market efficiency becomes solely down to arbitragers. As Gatev et al. (2006) point out, arbitrage profits are compensation for enforcing a more efficient market. Therefore, the profits made from exploiting price deviations across exchanges are compensation for providing a more efficient market within cryptocurrencies.

In the following report, we provide several new findings in the cryptocurrency trading and market

[1]On January 9 Bitcoin was 43% higher in South Korea than it was in the US a phenomenon known as Kimchi premium

efficiency literature. We first present our findings before we provide a possible explanation for these facts. For this analysis, we use data at the 10-minute frequency for 5 exchanges across Europe and the US. First, we find that between 1 April and 14 August the dynamic of price deviations across exchanges change. The start of the sample period experiences low, stable spreads between exchanges; however, towards the end of the sample, we see much larger recurring price differences. Our findings suggest that during this time period, Bitcoin markets become less efficient. This contradicts the majority of the Bitcoin market efficiency literature that suggests Bitcoin is inefficient but moving towards efficiency. Secondly, we find that a traditional pairs trading strategy performs the best to exploit these deviations in prices across exchanges. We follow the methodology of Gatev et al. (2006) using a formation period followed by a training period. We use the formation period to decipher the optimal trading strategy. We analyse both a traditional pairs trading strategy and those strategies put forward by cryptocurrency literature. We find that the strategy obtains profits larger than those reported by Alexander and Heck (2019b), suggesting that the spreads across exchanges have diverged. Finally, we find that realized volatility plays a big part in determining the spreads and that in most cases the use of volume in explaining the spreads is statistically insignificant. We calculate daily-realized volatility and aggregate spreads to a daily level, and find that realized volatility in some cases explains up to 80% in the variation of daily spreads.

The rest of the report goes as follows: Section 2 provides an analysis of related literature. Section 3 describes the data employed and provides the procedures followed to obtain and clean the data. Section 4 gives a brief descriptive summary of Bitcoin and the spreads across exchanges. Following this, section 5 gives an overview of the methodology employed to gauge our results. Section 6 provides the results of the optimisation within the formation period. Section 7 gives an analysis of the trading period and the strategy employed. Section 8 investigates the role that volume and volatility play in causing the spreads across exchanges to diverge. Finally, section 9 summarises and concludes.

2 Literature Review

This research project has close links to several streams of literature, particularly the cryptocurrency, statistical arbitrage, and market efficiency literature. The following section outlines and analyses the proceeding research that relates to the goals of this project. First, analysing the previous cryptocurrency market efficiency literature; following this; we give a more in-depth analysis of those papers that employ trading strategies to exploit exchange arbitrage. Finally, we will discuss pairs trading literature.

Bitcoin market efficiency has had quite a lot of exposure within academic literature. The prevailing opinion is that Bitcoin, alike to many markets, does not entirely constitute a fully efficient market according to the efficient market hypothesis. However, most research has suggested that over time, Bitcoin is becoming more efficient with the introduction of CBOE and CME futures furthering that. Fink and Johann (2014) were one of the first to study Bitcoin market efficiency and market microstructure. Their research concentrates on the market prior to 2014. The paper does not point towards an efficient market instead reporting that Bitcoin is highly volatile, exhibits high returns, is not informationally efficient and prices between exchanges are cointegrated. Fink and Johann were the first to report these stylised facts about Bitcoin, which are commonly known to those interested in cryptocurrency markets. Research preceding this formally tests for market efficiency within the context of the efficient market hypothesis. Urquhart (2016) implements a range of contrasting procedures to test for market efficiency on two non-overlapping sample periods: one between 2010-2013 and another from 2013-2016. The paper finds that the second sub-sample is more efficient than the first, suggesting that Bitcoin is inefficient but moving towards an efficient market. Nadarajah and Chu (2017) study succeeds Urquhart (2016). Their research shows that a simple power transformation leads to the satisfaction of the efficient market hypothesis. Nadarajah and Chu found a weakly efficient Bitcoin when testing the efficient market hypothesis using an odd-numbered power on Bitcoin returns.

More recently, research testing Bitcoin market efficiency does not merely point toward an efficient market with differing prevailing opinions. Khuntia and Pattanayak (2018) used the adaptive market hypothesis framework - developed by Lo (2005) - to test Bitcoin efficiency over time. Following the adaptive market framework, their study finds periods of efficiency and inefficiency within Bitcoin markets. Events such as the closing of Mt.Gox caused periods of inefficiency and the shifts in policy in favour of Bitcoin cause market efficiency. Overall, the paper finds evolving efficiency in Bitcoin markets, suggesting that significant events and structural changes cause efficiency to evolve. Contrastingly, Zargar and Kumar (2019) study Bitcoin information efficiency at the 15-minute level and find that there is a presence of information inefficiency at high frequencies. Zagar and Kumar follow the adaptive market framework and find that it points to the randomness of the Bitcoin Market at the daily level. Other research has concentrated on the role exchanges play in market efficiency. Borri and Shakhnov (2018) study the efficiency of cryptocurrency markets, finding that the discounts or premiums on exchanges are due to liquidity and counterparty risk. We intend to extend their research by examining the role of volatility and volume on the discounts of premiums on exchanges.

Bitcoin research is not solely made up of market efficiency literature. Urquhart (2018) test for the causes of investors attentions within Bitcoin. Urquhart employs Google Trends data and finds that realized volatility and volume are both significant drivers of the next day attention of Bitcoin. However, they find that attention offers no significant predictive power for realized volatility or returns. This is of interest to us, as we examine the relationship between realized volatility, volume and price deviations across exchanges.

There has been limited research that formally tests for Arbitrage profits with cryptocurrency markets; however, Makarov and Schoar (2019) and Alexander and Heck (2019b) employ statistical tests and trading strategies to both exploit and explore deviations in prices across exchanges. Both the papers relate closely to the research objectives of this paper. Makarov and Schoar (2019) set out to investigate exchange arbitrage within the cryptocurrency space, concentrating on the deviations of Bitcoin price across regions. While the paper does not set out with a clear hypothesis, it does propose that the deviations across regions are likely due to market frictions, the unregulated nature of cryptocurrencies and the investor's attitudes within each region. The paper finds that those exchanges outside of Europe and the US react greater to positive sentiment regarding Bitcoin, leading to higher deviations across regions than within. Although the paper does not investigate causes for deviations within regions, the analysis does, however, start to address the reasons why prices diverge. Finally, Makarov and Schoar find that between November 2017 and February 2018, the available arbitrage profits are around $980 million. However, this figure has little significance in describing trading profits, as it does not consider market frictions.

Alexander and Heck (2019b) is the first paper to implement and record the possible trading profits from exchange arbitrage strategies. The paper sets out to examine the role of the CBE and CMOE futures in price manipulation of Bitcoin, price discovery and examining the market efficiency of Bitcoin, with the hypothesis that if futures play a leading role in price discovery, then manipulative trading on those contracts could influence the spot price of Bitcoin. The reason for this hypothesis is that the Bitcoin futures contracts are too large for hedging, leaving the demand for these to be almost wholly made up of speculation. The paper finds that during periods of extreme price declines and the most recent bull market, there were indications of manipulation in both CBOE and CME contracts. Finally, Alexander and Heck implement a simple pairs trading strategy following the works of Makarov and Schoar to find whether there are significant profits from arbitrage strategies. The paper uses data at a 10-minute frequency for the five leading exchanges. They find that between April 2018 and the end of March 2019 - when including fees - trading profits are insignificant, finding that these hedge fund strategies have been successful in increasing market efficiency. Furthermore, Alexander and Heck (2019a) recently released another paper further investigating the role of CBOE and CME futures within price discovery. The paper finds that during our sample period, there were two instances of clear manipulation. Firstly, between 12 and 15 May 2019 when prices rose 15% and between the 25 and 27 June 2019 when the spot price increased by 10%, there are strong contributions of price discovery from CME futures.

Traditional pairs trading has had extensive research; however, there does not appear to be a consensus regarding its most profitable market conditions. Older literature suggests that the strategy is market neutral; however, more recent studies have found that it is most profitable in periods of turmoil. Gatev et al. (2006) paper are one of the first and most influential papers to introduce the popular hedge fund strategy of pairs trading and find that the strategy is market neutral. The paper sets out to investigate the profitability of the trading strategy for US stocks from 1962 to 2002 using data at a daily frequency. However, it is noteworthy that the paper does not use robust testing for pairs, neglecting the cointegration method and using the sum of squared errors. The paper uses a formation period where they optimise and pick pairs for 12 months and trade using these portfolios for the following six months. The study finds an average annualised excess return of 11% with a low level of risk and follows the rhetoric that the profits are compensation for enforcing the law of one price. Bowen and Hutchinson (2016) implement the pairs trading strategy in the UK equity market. The paper implements the trading strategy between 1980 to 2012, finding an average annual return in the range of 6.2% to 9.6%. However, most interestingly, the paper finds that the pairs trading strategy provides abnormal returns during a crisis, with the portfolios delivering returns of 36% to 48% over two years during the 2008 financial crisis. Succeeding this Do and Faff (2010) extend Gatev et al. (2006) research by extending the sample to the beginning of 2008. The paper finds that the trading profits are diminishing as more of the pairs converge. Finally, Perlin (2009) implement the trading strategy in Brazilian financial markets finding the strategy to be market neutral.

3 Data Description

For this research project, the spot price of Bitcoin for five different exchanges at the 10-minute frequency was obtained, alongside the volume of each trade. Within this section of the paper, we discuss the reasons why we choose the data, the procedures followed in cleaning the data and the summary statistics.

Firstly, before we go into the process of obtaining the data and cleaning it, we believe it is paramount to point out the findings of Alexander and Dakos (2019). The paper draws attention to the reliability of many sources of cryptocurrency data and research using non-traded data. The paper points out that a large proportion of the research projects published within the subjects of cryptocurrency portfolio optimisation, trading development or hedge fund analysis do not use traded data. This can often lead to incorrect findings when it comes to empirical work. The paper intends to be a guide to empirical research of cryptocurrency; therefore, we follow the advice of the paper to make sure we have a robust source of data. The paper suggests that one of the most reliable traded data sources is Coinapi, which allows the retrieval of historical order books using data transfer protocols, known as API protocols (Application programme interfaces)[2]. Therefore, we use Coinapi's to obtain the data for this research project.

To analyse exchange arbitrage, we obtain data from five exchanges; Coinbase, Bitstamp, Kraken, Gemini and ItBit. Here, we outline the reasons why we select these exchanges. Firstly, the analysis of Alexander, C and Heck, D looks at these exchanges, and we intend to follow on from their research. Secondly, Heck provides us with the data he used for his analysis, giving us a 'practise' set of data, to make sure our code implements the trading strategy correctly and gives the results that he reported, making it easier compare results. Thirdly, as Bitwise Asset Management points

[2]Coinapi is a paid service which provides API's to 204 exchanges and 6483 assets and allows a small quantity data to be downloaded with a free 'key'

out, these five exchanges do not inflate their trading volume with wash trading; therefore, we can be sure that we could trade at the prices and volume reported. Finally, the exchanges we look at include the three most popular exchanges in both America and Europe, Coinbase, Bitstamp and Kraken, as well as two smaller exchanges with ItBit and Gemini. The variability in the sizes of the exchanges may give an insight into how sensitive the spreads are on trading volume as well as other dependent factors. In the following section, we will give further details on each exchange.

Using Coinapi's URL calls, we obtain data at the 10-minute frequency for the sample of the 1 April 2019 to 14 August 2019; this gives 19,437 inputs for each of the five exchanges. We choose this sample period for several reasons. Firstly, it is the most up to date data we can use without infringing on the completion of the empirical study. Secondly, we believe that it is during this period that deviations across exchanges begin to widen. Finally, Alexander and Heck's sample period finishes on the 31 April 2019; therefore, our work picks up where they left off. We acquire data at the 10-minute frequency as it is the highest frequency that we can obtain reasonably. Coinapi's has a limited number of calls for a free key; therefore, even at this frequency, we create multiple email addresses and get family and friends to request free keys. Once we obtain the data, the process of cleaning it involves adding time-steps where there is missing data, made more complicated as excel does not recognise the time-stamp. This can be worked around by separating the time and date column and then adding them back together. The larger exchanges had 2 to 3 missing time steps, whereas the smaller ones had a much more significant proportion of the data missing. For trading analysis, once a pair is put together if there is a missing value for either exchange we delete the entire row. This way it does not interrupt the moving averages used in the trading strategy.

4 Statistical Summary

Within this section, we provide a statistical summary of the data used within the project. First, we present a statistical summary of Bitcoin and then follow on with spreads between exchanges. For the analysis of Bitcoin, we create a Bitcoin price index. We create the price index by multiplying each exchange by their volume taking the sum of all the exchanges and then dividing by the total volume. Therefore, exchanges with larger volumes will play a larger part in determining the index price. Table 1 provides; the mean, standard deviation, minimum and maximum, skewness and kurtosis. Bitcoin is incredibly volatile, negatively skewed and highly leptokurtic (fat-tailed) with a kurtosis above the normal distribution. The sample exhibits a minimum price of \$4076.96 and maximum of \$13849.80.

Figure 1 shows the 10-minute interval of Bitcoin index price and Figure 2 shows the annualised 10-minute interval volatility of Bitcoin. The price path includes a significant increase in price, starting at \$4091.12 and reaching heights of \$13849.81 at the end of June. Following the heights of June, the price drops off before beginning to recover again at the start of August. Our sample period coincides with the two periods where Alexander and Heck conclude that CME futures had a leading role in spot price discovery. Between 12 and 15 May we saw an increase of 15% and between 25 and 27 June we see a 10% increase lead by CME futures price discovery. This tranquil period is shown in both the spot price figure and ExWMA. We capture the volatility of Bitcoin throughout the sample period using an ExWMA of order 0.99 as we believe this captures volatility best at such a high frequency. During the sample period, the highest points of volatility are those periods where CME futures drive the spot market.

Figure 1: Bitcoin Price Path

Figure 1 displays the price path of Bitcoin index between the 1 April 2019 to 14 August 2019

Figure 2: Bitcoin Volatility

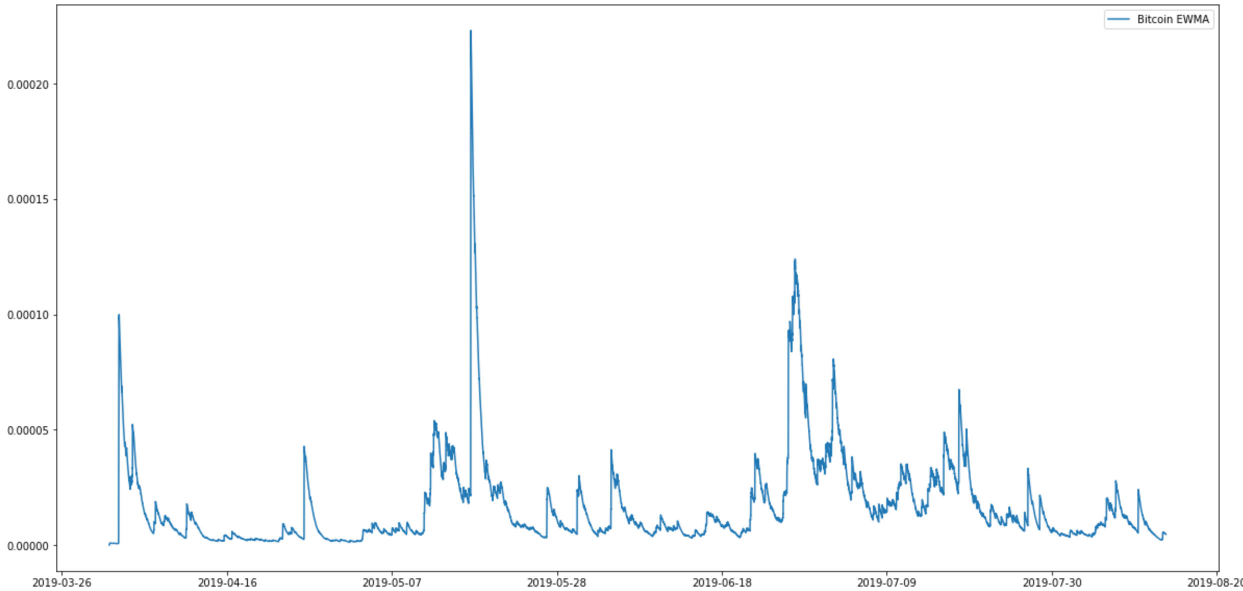

Figure 2 displays the volatility of the Bitcoin index using a λ =0 .99 between the 1 April 2019 to 14 August 2019

Correspondingly, the volume of Bitcoin traded on each exchange coincides with the periods of high volatility and CME price discovery. Figure 3 depicts the volume of each exchange throughout the sample period. Between 12 and 15 May we see large trade volume as well as between 25 and 27 June. Finally, the graph shows that Coinbase has the largest trading volume, followed by Bitstamp and Kraken, with ItBit and Gemini experiencing the lowest trading volumes.

Figure 3: Bitcoin Volume

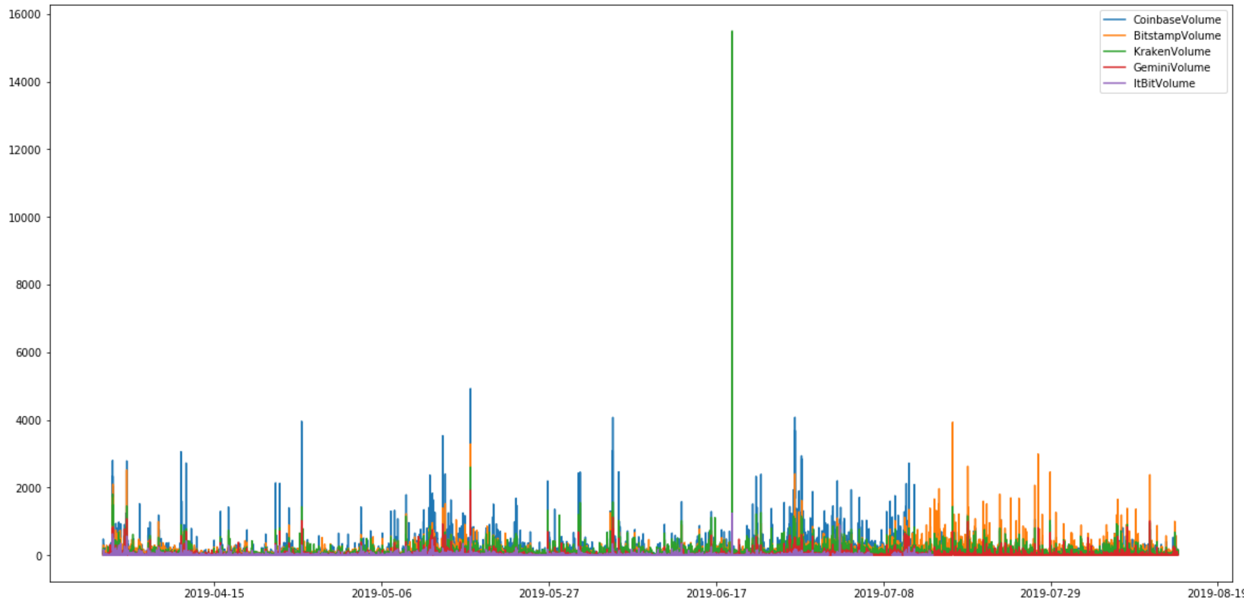

Figure 3 displays the volume of Bitcoin traded for each of the five exchanges between the 1 April 2019 to 14 August 2019

Within this research project, we will mainly deal with the 'spread' or difference in the price of Bitcoin between each exchange; for that reason, we provide a descriptive summary of the spreads. The spread is calculated as the difference in prices between each exchange as a percentage of the exchange, that is the basis of the calculation. As the spread has a different basis of the equation, they are not symmetric, for example, a Coinbase Bitstamp spread is not the same as a Bitstamp Coinbase spread. Figure 4 depicts the spreads between the prices of several exchanges from 1 April 2019 to 14 August 2019. There is a distinct change in the behaviour of the spreads within the sample. The second half of the sample sees the spreads widen, in comparison to the start of the sample, where we see low spreads and low volatility. Interestingly, the periods where CME futures lead price discovery witnesses short-lived large increases in spreads. Throughout June we see large recurring price deviations across exchanges which appears to be led by spreads including Coinbase.

Figure 4: Price Spread Between Exchanges

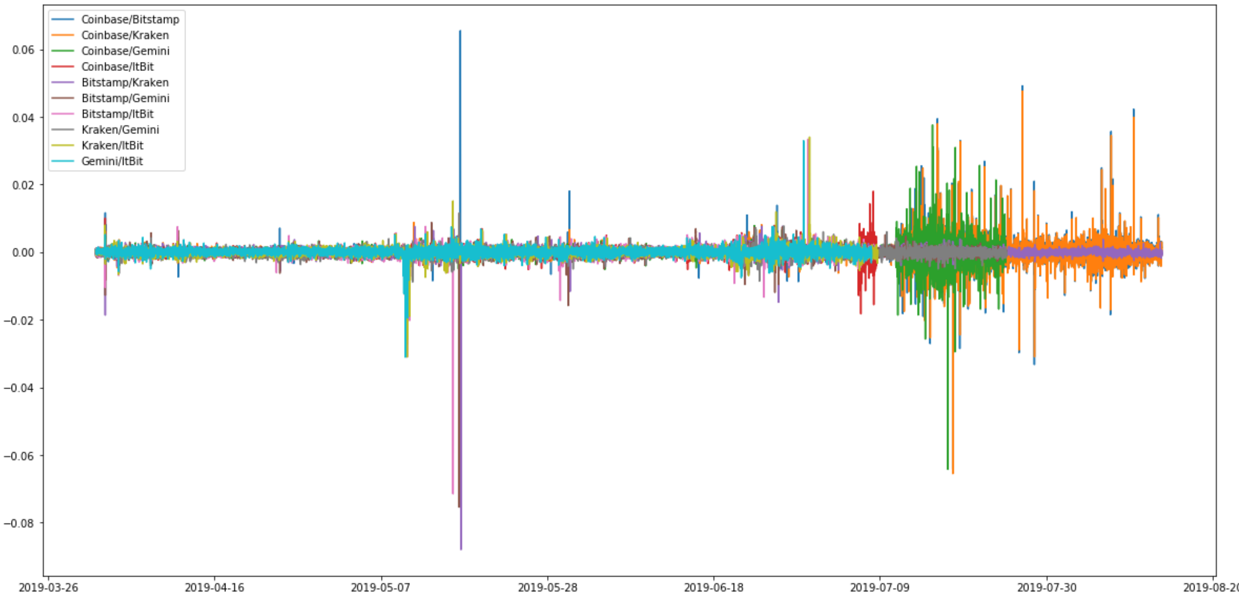

Figure 4 displays the price spread of Bitcoin between each exchnage between the 1 April 2019 to 14 August 2019

Furthermore, we compute some summary statistics for the spreads in matrix form, provided in Table 3. The exchange in the row is the basis of the calculation. The largest average spread is between Coinbase and Kraken, this is quite surprising as these exchanges have the largest volumes, but Kraken is on average 8.12 basis points cheaper than Coinbase. Bitstamp appears to be trading at a discount in comparison to the other exchanges. The rest, however, do not appear to follow any clear pattern. Coinbase and Bitstamp had the most volatile spread, with Coinbase experiencing the average largest overall standard deviation. Bitstamp and Kraken had the largest; spread, kurtosis and skewness. The large kurtosis tells us there are outliers of a significant degree within the spreads, which potentially could provide large profits. Bitstamp unsurprisingly had a negative skew throughout, as it trades at a discount throughout.

5 Exchanges and Fees

Before we move on to the methodology, we think it is useful to cover some necessary information about the exchanges included. As mentioned in the data section, we are looking at the American exchanges Coinbase, Kraken, ItBit, and the London based Bitstamp. All these exchanges allow the trade of fiat to multiple cryptocurrencies and include trading from both the dollar and pound. Of the five, only Kraken allows for short sales. However, Bitstamp has hinted that it may come in the future Bitstamp-Lucas and Reddit.com[3]. The business model of the exchanges works based on charging a small fee or small percentage of each trade. The percentage is based on the trader's volume throughout the month; traders are rewarded for larger volumes with lower fees. We have presented the fees structure for each exchange in Table 2. As shown, they range from 0.35 to 0.16 depending on trading volume. Finally, it is noteworthy that these exchanges make up half of the ten exchanges, which Bitwise asset management suggest, do not report inflated volume.

6 Methodology

Within this section, we outline the trading strategies employed within the project. Following Gatev et al. (2006), we use a formation period or an in-sample period to optimise the trading strategy and then implement it in a trading period or out of sample period. Therefore, we split our sample period in half, the training period starts on the 1 April 2019 and ends on the 7 June 2019 with the full sample ending on the 14 August 2019. Within the training period, we test two trading strategies; first, the trading strategy put forward by Alexander and Heck (2019b), subsequently, we employ a more traditional pairs trading strategy suggested by Gatev et al. (2006). Following this, we optimise each strategy.

To perform an analysis of the arbitrage profits, we apply two pair-trading strategies. Heck and Alexander extended the analysis of Makarov and Schoar (2019). The paper suggests a trading strategy of holding Bitcoin and fiat in wallets of all the exchanges, and then when the spread is large enough to cover trading costs, buy the winner and short the loser, then holding this position until the trade is reversed, profiting from mean reversion. To implement this strategy we calculate the percentage spread between synchronous prices on the exchanges and divide this by the basis exchange, giving us a spread in percentage terms. Thereafter, we calculate the moving average and the moving standard deviation of the spread. An upper Bollinger Band is formed as the MA plus two multiples of the MSD, and a lower Bollinger Band as MA minus two multiples of the MSD. If the spread goes beyond the upper Bollinger Band after deducting fees, we sell or short the spread by selling one Bitcoin on the exchange that is the basis of the equation and buying on the second

[3]In 2018, Binance answered a question regarding whether short selling would be an available feature in the future. They suggested that it is something they are working on bringing it to the site.

exchange. If the spread goes beyond the lower Bollinger Band after deducting fees, we buy the spread, by buying one Bitcoin on the exchange that is the basis of the equation and sell one Bitcoin on the second exchange. The trade is released when the spread reverses. For capital controls and risk-management reasons, we only open one trade between a pair at a time. This is because it limits the amount of capital open to divergence risk and means that we only have to hold one Bitcoin and fiat on each exchange. We present an example of the trades in Figure 5 as well as all the equations needed below.

$$S_i = \frac{Exchange1 - Exchange2}{Exchange1} \tag{1}$$

$$BolingerBand_i = MA_i \pm 2 * MSD_i \tag{2}$$

$$MovingAverage = u_{n,t_0} \tag{3}$$

$$MovingVariance = \sigma^2_{n,t_0} \tag{4}$$

$$u_{n,t_0} = \frac{1}{n} \sum_{i=1}^{n} Y_{t_0-i} \tag{5}$$

$$\sigma^2_{n,t_0} = \frac{1}{n-1} \sum_{i=1}^{n} (Y_{t_0-i} - u_{n,t_0})^2 \tag{6}$$

$$u_{\lambda,t_0} = (1-\lambda) \sum_{i=1}^{n} \lambda^{i-1} Y_{t_0-i} \tag{7}$$

$$\sigma^2_{\lambda,t_0} = (1-\lambda) \sum_{i=1}^{n} \lambda^{i-1} (Y_{t_0-i} - u_{\lambda,t_0}) \tag{8}$$

We sell one BTC on exchange i and buy one BTC on exchange j if:

$$\S_i - \tau_{i,j} > max(MA_{i,j} + 2MSD_{i,j}) \tag{9}$$

We buy one BTC on exchange i and sell one BTC on exchange j if:

$$\S_i + \tau_{i,j} < min(MA_{i,j} - 2MSD_{i,j}) \tag{10}$$

where $\tau_{i,j} = \rho \frac{S_i + S_j}{S_i}$ are the trade costs as a percentage of the basis exchange and ρ is the transaction fee.

Figure 5: Trading Strategy

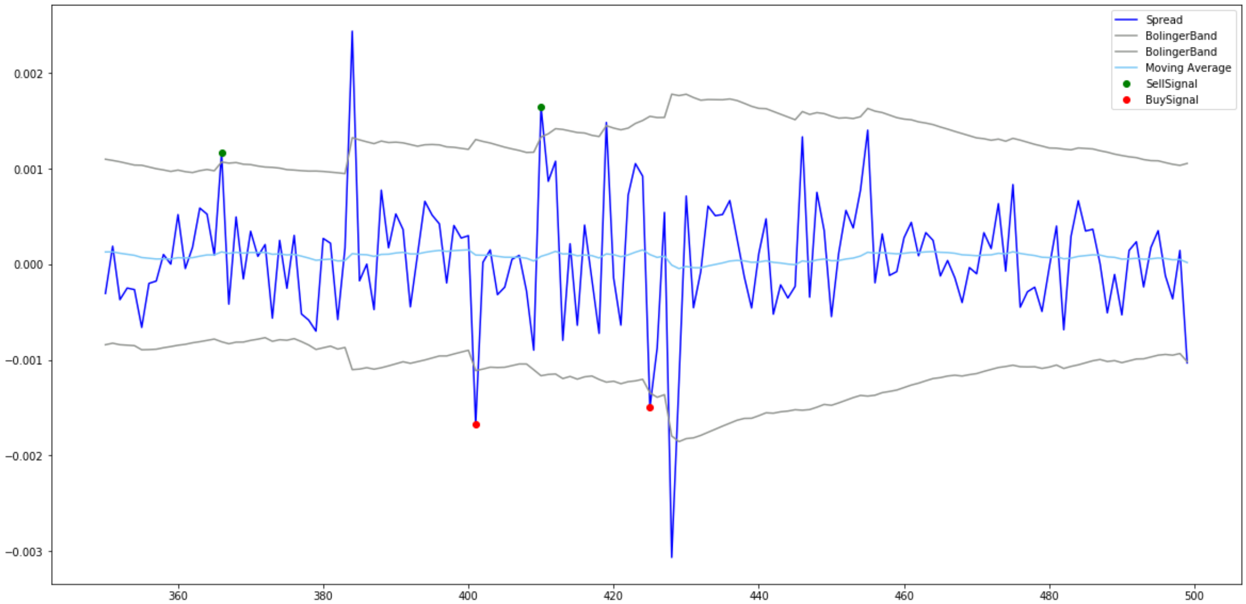

Figure 5 demonstrates the trading strategy employed. The green dot denotes a buy signal and a red denotes a sell signal

Second, we employ a traditional pairs trading strategy. Again we open a trade when the spread goes beyond the Bollinger Bands. However, the strategy differs as we release the position once the spread crosses the moving average. So, the spread goes beyond the Bollinger Bands, and once it reverts to its mean (crosses its moving average), we release the position. Therefore, we profit from mean reversion of the spread. Traditional pairs trading was first implemented to exploit two cointegrated stocks with the idea that if they diverge from this the cheaper stock will increase in price and the more expensive stock will decrease in price or revert to their mean. However, we believe that the strategy will successfully exploit exchange arbitrage spreads. The implementation of this trading strategy will likely see trades open and close at a higher frequency; however, we suspect it is more likely to trade at a loss.

Between the two trading strategies, there are then three interchangeable parameters. Firstly, the moving average parameter. The moving averages can either be an equally weighted moving average (MA) and moving standard deviation (MSD) of some order n, or an exponentially MA and MSD of some order of the smoothing parameter lambda. The equally weighted moving average will account for all spreads within the window, whereas the exponentially weighted will account for all but reduce the weighting exponentially as time passes on. Secondly, the order of window size within the moving average and the order of lambda within the MA is adjustable. A larger window length or lambda will see a smoother Bollinger Band whereas a shorter n will see a spikier Bollinger Band. Finally, the order of multiple of the MSD included in the Bollinger Band is adjustable. A lower order of multiple will mean trades open at a more frequent rate. The advantage of trading at a higher multiple of MSD means that when trades open, it should make more significant profits.

To optimise the trading strategy, we consider the arbitrage profits for the different parameters for each trading strategy. A hedge firm that deals explicitly with cryptocurrencies is the most likely candidate to employ the trading strategy; therefore, we assume a high trading volume (somewhere between $100,000 and $1,000,000). This would lead to a trading fee of around 0.002 or 20 basis points. For simplicity, we assume a trading fee of 20 basis points on all exchanges. We are aware that it is not realistic to be paying identical fees on each exchange; however, we believe this will not cause bias within the results. Figures 6 7 8 9 10 and 11 show arbitrage profits from all the spreads resulting from different smoothing parameters for ExWMA, different window lengths for EqWMA, and different multiples of MSD added to the Bollinger Band. We have implemented a window length between 5-25, a lambda from 0.8-0.99 as well as multiples of 1 and 2 for the MSD.

7 Optimisation Results

In the following section, we review the performance of each trading strategy within the training period. To start, we demonstrate how the optimisation works before considering each trading strategy, and finally introduce fees. The section will go as follows: first, we consider each strategy separately finding the optimal parameters. Following that, we compare all optimised strategies while considering fees. We include an analysis of the equally weighted moving average (EqWMA) and the exponentially weighted moving average (ExWMA) strategy for both the first and second trading strategy. Finally, we will summarise all of the trading strategies and pick an overall optimal strategy to implement in our trading period.

First, we demonstrate how the optimisation works. As mentioned in the previous section, each trading strategy includes the construction of a Bollinger Band. We construct the Bollinger Bands with either an exponentially weighted MA and MSD or an equally weighted MA and MSD. The moving averages and MSDs are dependent on the order of lambda or window length. To optimise this, we created a function on python for each trading strategy with window length or lambda as changeable parameters and the profits of the strategy as the output. We then run the function through a for loop with the lambda or window length as the index. The function creates a new variable with the different profits for each order of lambda or window size. Figure 6 demonstrates this.

We now investigate how the arbitrage profits vary from differing parameter, assuming zero trading fees. We now consider the first strategy using ExWMA and EqWMA with two multiples of MSD. Figure 6 and 7 depict the arbitrage profits from the exchanges considered, resulting from an ExWMA smoothing parameter between 0.8 and 0.99, and an EqWMA order between 5 and 30. Figure 6 and 7 shows a robust result where all pairs appear to react the same to differing EqWMA orders and smoothing parameters. It appears that the strategies performance increases with an increase in EqWMA order or ExWMA smoothing parameter. The equal-weighted strategy sees an increase similar to a logistic model, whereas the exponentially weighted portfolio sees a more a gradual increase. A smoothing parameter of 0.98 results in the largest profits for the ExWMA strategy, generating $47,801. An EqWMA order of 26 generates the largest profit for the equally weighted strategy, generating $55,875.36. We now do the same analysis using a Bollinger Band constructed using just one multiple of MSD, presented in Figure 8 and 9. First, we notice that the strategy generates a much larger profit per pair. Moreover, we see a differing reaction to the previous optimisation for the changing order of EqWMA and smoothing parameter. The plots show an almost 180-degree line with the profits staying equal until we approach the highest order considered where the profits start to diminish. EqWMA order of 6 performs the best-generating profits of $191,490.70, and an ExWMA smoothing parameter of 0.845 generates profits of $181,477.90.

Figure 6: ExWMA Optimiser

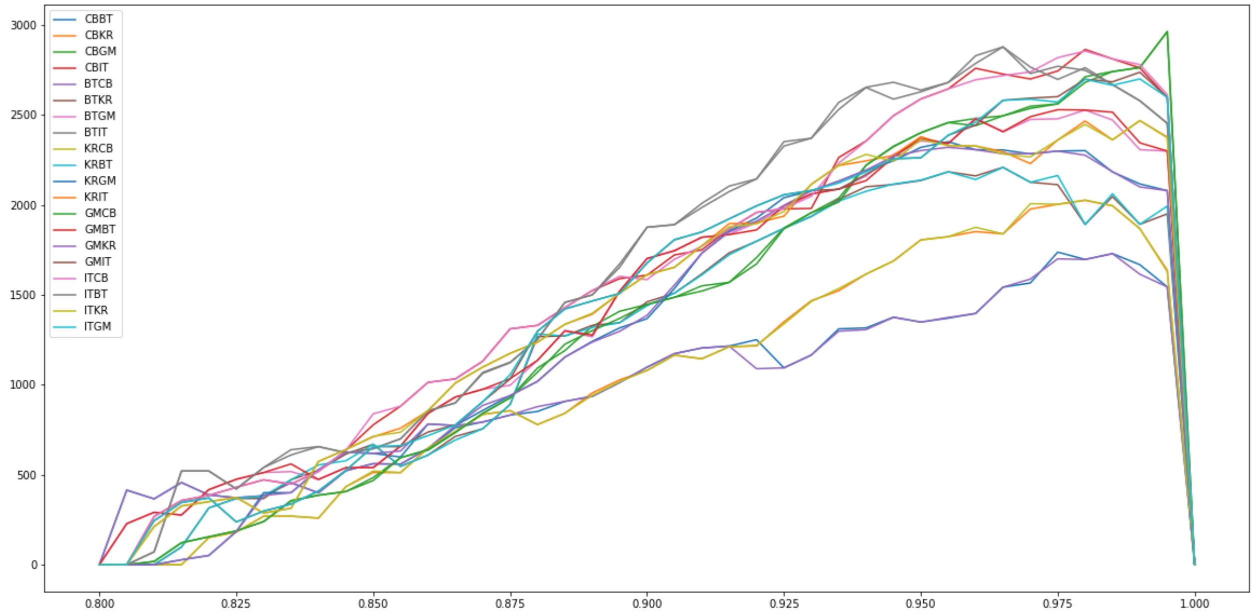

Figure 6 depicts the trading profits from differing orders of λ for the first strategy assuming 2 multiples of MSD. l

Figure 7: EqWMA Optimiser

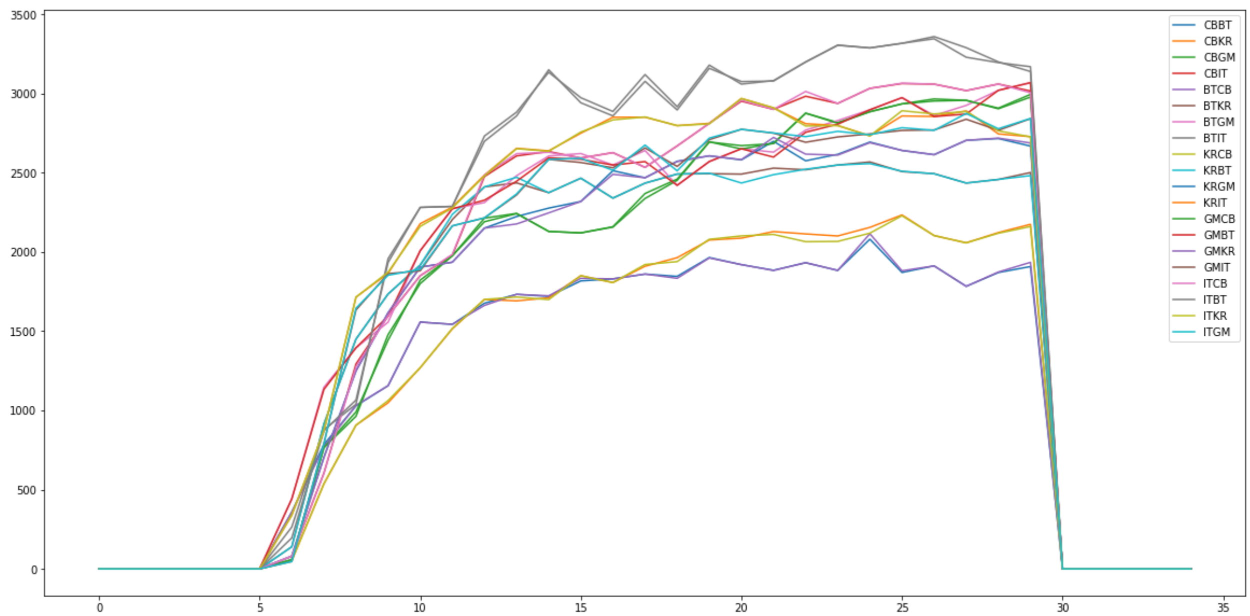

Figure 7 depicts the trading profits from differing orders of EqWMA for the first strategy assuming 2 multiples of MSD. l

Figure 8: ExWMA Optimiser constructed with 1 MSD

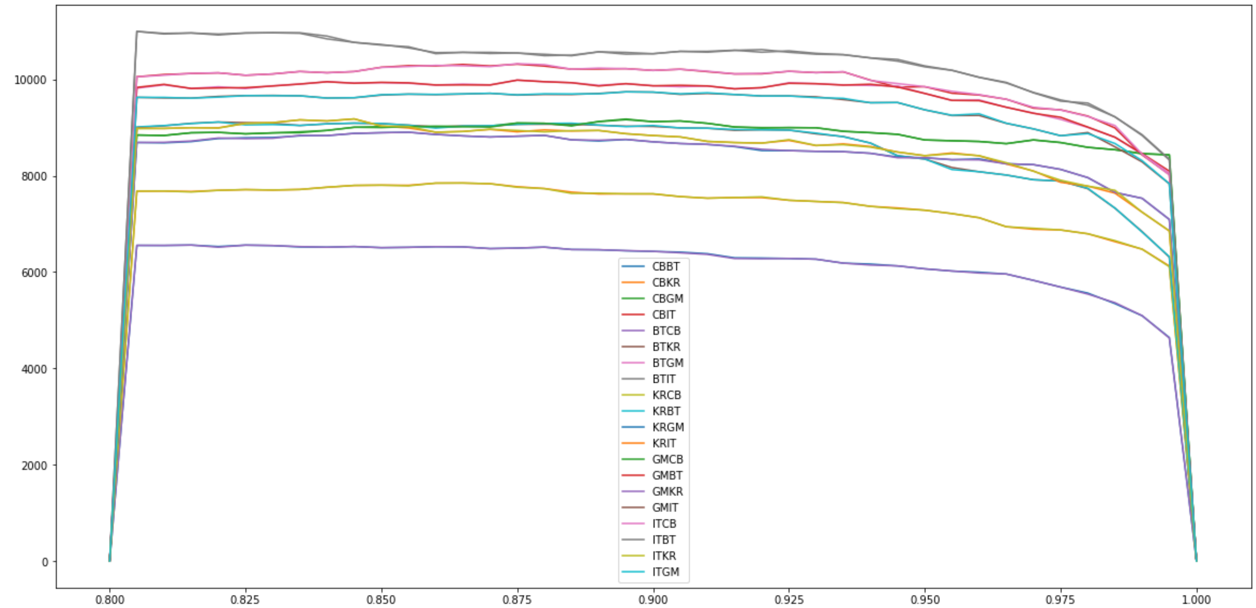

Figure 8 depicts the trading profits from differing orders of ExWMA for the first strategy assuming 1 multiple of MSD.

Figure 9: EqWMA Optimiser constructed with 1 MSD

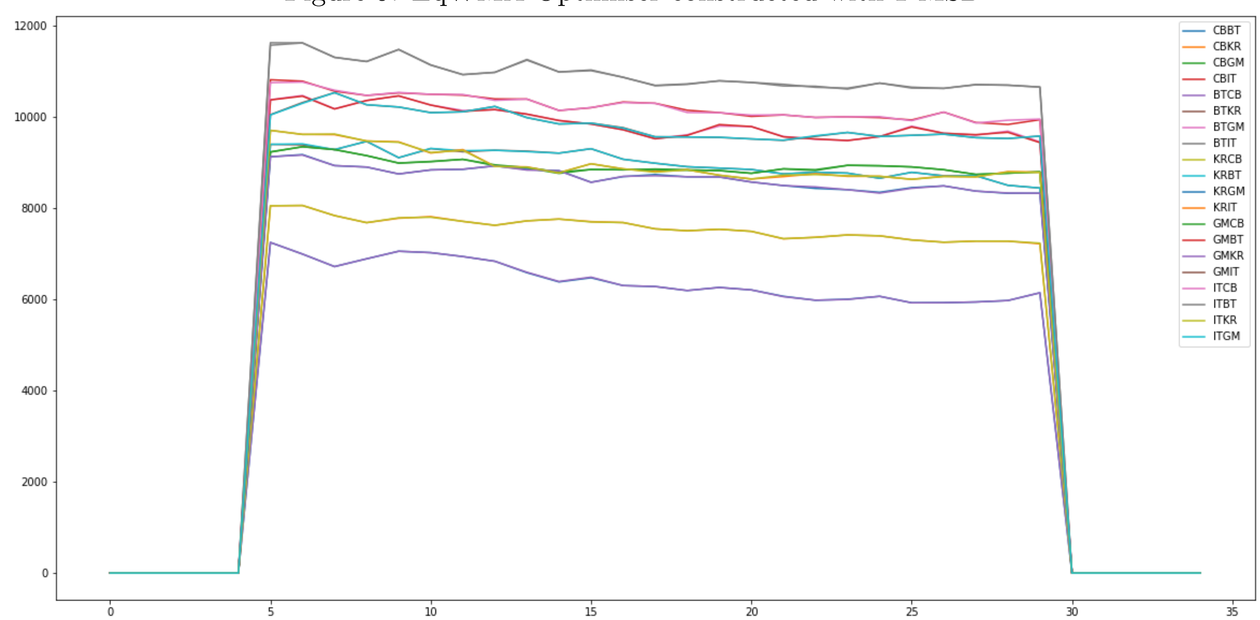

Figure 9 depicts the trading profits from differing orders of EqWMA for the first strategy assuming 1 multiple of MSD.

Finally, we look at a traditional pairs trading strategy put forward by Gatev et al. (2006). For this trading strategy, we only include the analysis of a Bollinger Band constructed with two multiples of the moving standard deviation. Each trade must account for both the opening fees and closing fees; therefore, it is more likely to incur a loss. Hence, we only implement two multiples of the MSD to safeguard against this. In comparison to the other trading strategies, trades should open at a higher frequency as they close at a faster rate. Figure 10 plots the different arbitrage profits from different orders of λ. Similarly, to the other trading strategies, the plot shows a robust relationship between profits and order of lambda, with a higher λ resulting in higher profit. A λ =0.98 obtained the most profits at \$34358.86 with only one pair not obtaining the most profits at that order of λ. Furthermore, we optimise the strategy using the equally weighted constructed Bollinger Band. Figure 11 plots the different arbitrage profits from different orders of EqWMA. The plot shows the profits increasing rapidly at the start; however, this appears to diminish quite quickly before levelling off. Following the same rhetoric, the plot shows a robust relationship between the profits and order of window length, with a higher window length resulting in a higher profit. The best performing window length was 29 that obtained profits of \$71782.91. Using the optimal order of λ for each pair obtained profits of \$72305.

Figure 10: Traditional pairs trading optimiser constructed with EwWMA

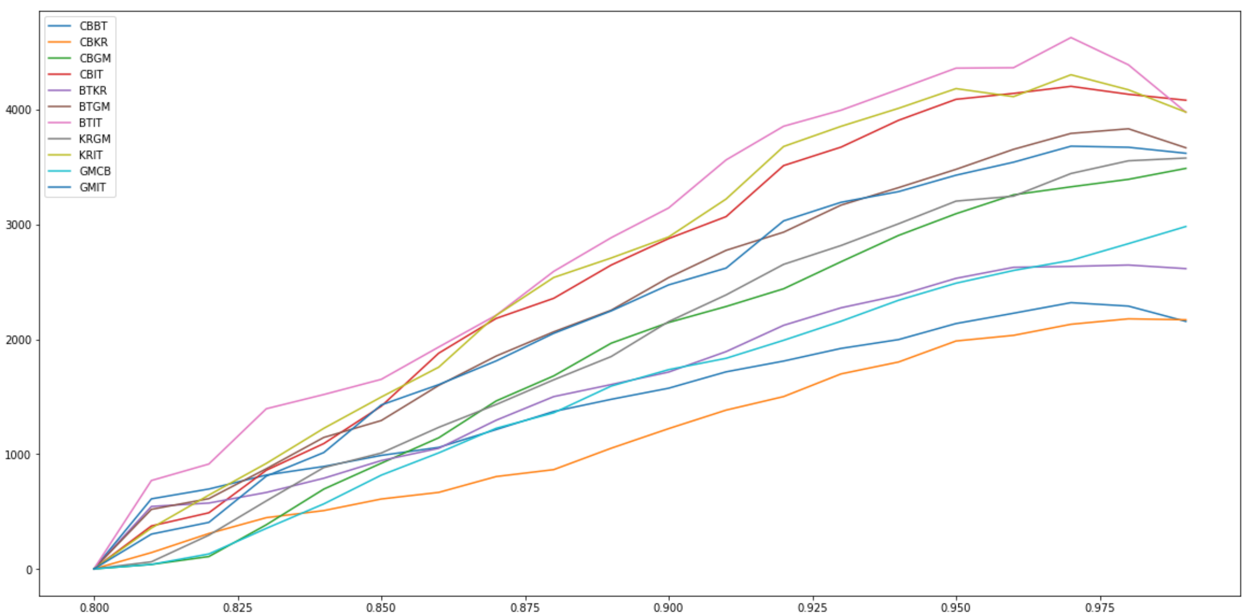

Figure 10 depicts the trading profits from the traditional pairs trading strategy for differing orders of λ.

Figure 11: Traditional pairs trading optimiser constructed with EqWMA

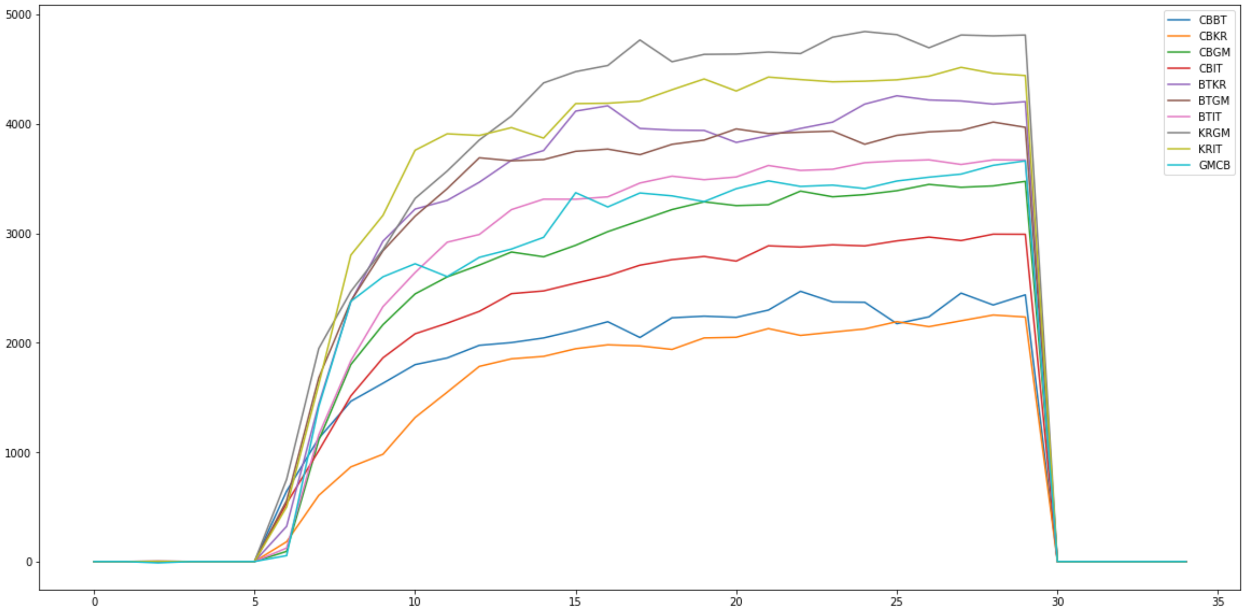

Figure 11 depicts the trading profits from the traditional pairs trading strategy for differing orders of EqWMA.

We now investigate how the arbitrage profits differ from each trading strategy when assuming trading fees. To see how they perform Table 4 shows the profits obtained in the training period. Based on the previous analysis, we consider an ExWMA smoothing parameter of 0.98 using a Bollinger Band constructed with two multiples of MSD and an order of 0.84 using a Bollinger Band constructed with one multiple of MSD. The ExWMA constructed with two multiples of MSD performs the better of the two, obtaining profits of \$3122.846. Already the presence of transaction costs drives down the available profits from the strategy. We obtain only 6.5% of the profits of the worst-performing strategy from the previous analysis. We now consider an EqWMA order of 26 using a Bollinger Band constructed with two multiples of MSD and an order of 6 using a Bollinger Band constructed with one multiple of MSD. The EqWMA constructed with two multiples of MSD performs the better of the two, obtaining \$2168.06. Finally, we consider both the EqWMA of order 29 of the traditional pairs trading strategy and the ExWMA with a smoothing parameter of 0.98 of the strategy. The EqWMA performs the better of the two and the best overall obtaining profits of \$3674.93.

Having considered all possible trading strategies and transaction costs, it is clear that the traditional pairs trading strategy using an EqWMA order 29 performs the best. It obtains \$500 more in profits than the second-best performing portfolio. The strategy made almost on average, ten more trades per pair. The portfolio appears to be more successful as it closes out the position at a much faster rate, meaning it can open up trades at a lot higher frequency. This trading strategy will be taken forward into the out of sample period with the hope that the strategy will be able to exploit a large percentage of the spreads seen in the second part of the sample.

8 Results

Having considered multiple trading strategies, we now implement the optimal trading strategies in the trading period. Table 5 presents the trading profits, and Figure 12 depicts the transactions for each pair. Overall the strategy outperforms the training period by 700% with trading profits of $25,725.133 and an average profit of $1286.25. Pairs including Coinbase performed the best with an average profit of $3175.24. The worst performing pair was ItBit and Gemini that incurred a loss of $29.11 and a profit of $2.65. Drawing our attention to the cash flow, notice that after each positive transaction, the trade incurs a loss. The initial opening of the trade profits from the spread, however, once the pairs converge, we close the position and incur a loss due to the trading costs. At the start of the trading period, we see a low frequency of trades that have a limited cash flow, however, in the second half, we see trades at a much higher frequency with higher cash flows. This suggests that we see larger profits towards the end of the sample. In comparison to Heck and Alexander's finding, the trading profits have seen a significant increase. Heck and Alexander test the profitability over a year with fees as low as ten basis points finding a total profit of $7829 - less than a third of the profits we obtain. Overall, the resulting profits contradict a large majority of the previous research. Alexander and Heck (2019b), Khuntia and Pattanayak (2018), Urquhart (2016) and Nadarajah and Chu (2017) all advocate for either efficiency or a market that is becoming more efficient, however the increase on arbitrage profits suggests that Bitcoin is becoming less efficient. This, does, however, support the findings of Zargar and Kumar (2019) who find the Bitcoin market is inefficient and Bowen and Hutchinson (2016) who find that pairs trading is more profitable in periods of high volatility (financial crises).

The resulting trading profits suggest that Coinbase has the largest deviations from the general price of Bitcoin. The prevailing research of Borri and Shakhnov (2018) would suggest that the cause of Coinbase's deviations would be either due to liquidity risk or counterparty risk, however, we suggest that it is unlikely due to two reasons. Firstly, as Figure 13 shows, Coinbase's distribution does not appear to skewed; therefore Coinbase, does not generally trade at discount or generally trade at premium. As it neither trades at a discount or a premium it hard to argue that these deviations are due to liquidity risk or counterparty risk. For those exchanges that do trade at a discount or premium, we would expect it to trade at a discount or premium most of the time as these risks on an exchange will only vary to a degree. Secondly, Coinbase is one of, if not the most reliable exchanges in terms of liquidity and counterpart risk, therefore, if these were the causes for the large spreads, Coinbase would only trade at a premium. However, this does leave us puzzled regarding why pairs, including Coinbase, were the most profitable. We attempt to investigate this in the following section.

Figure 12: Transactions of the trading period.

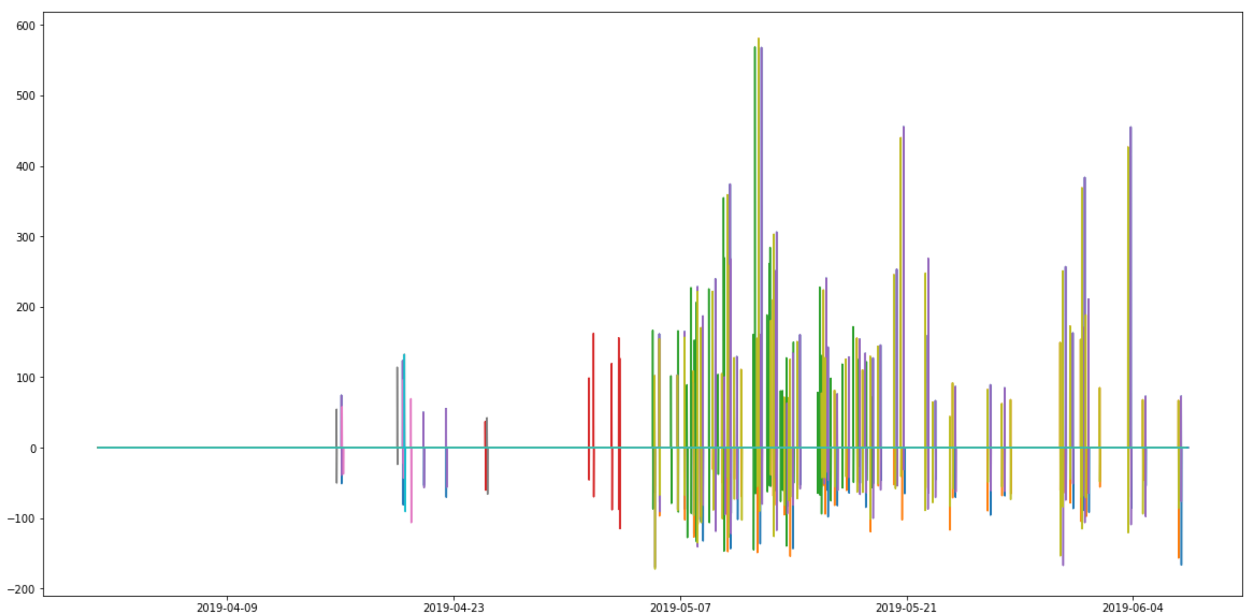

Figure 12 depicts the transactions, including all exchanges in our trading period. Given in USD

Figure 13: Histogram of Coinbase spreads.

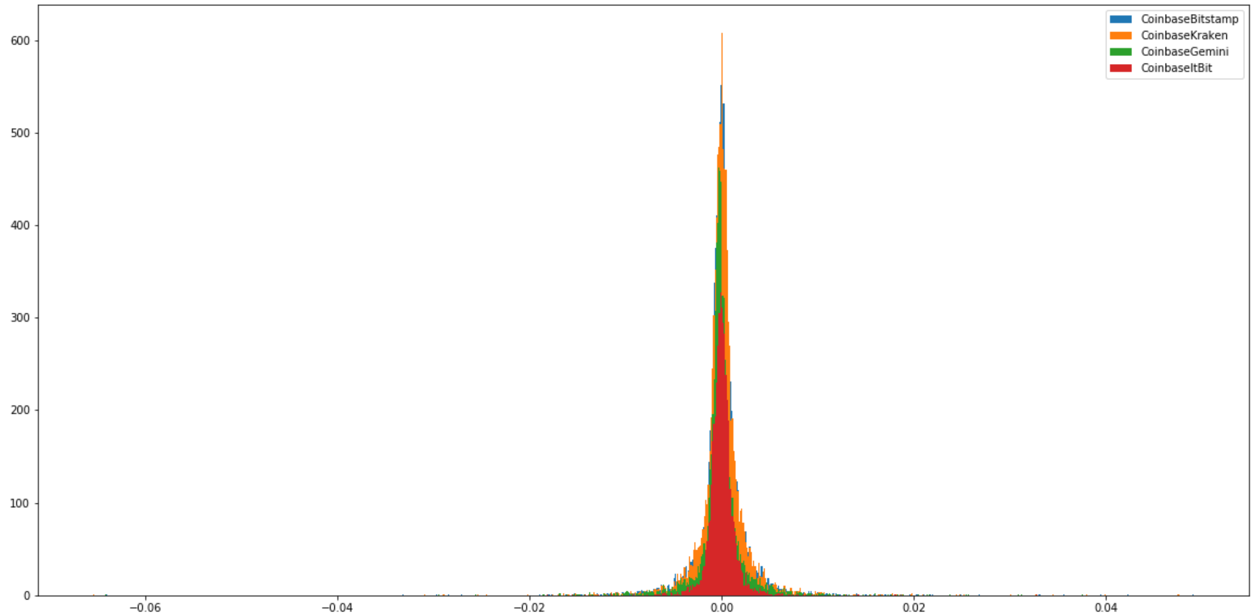

Figure 13 depicts the distributions of the deviations in price across exchanges for Coinbase, including all exchanges in our trading period. Given in USD

9 Volume and Volatility

We hypothesised that volatility would have some explanatory power over arbitrage spreads. We can begin to explain this through two rationales. Firstly, volatility means a tendency for prices to fluctuate; an increase in volatility may be perceived as an increase in risk; this can adverse investors away from the assets in question. The reduction in investors will lead to a less efficient market and rise in arbitrage profits. Secondly, if prices fluctuate and investors are unexposed to all available information there is a tendency for assets to be mispriced. In the context of this research, this could lead to deviations in the price of Bitcoin on different exchanges. To investigate this, we develop a regression to test the relationship between the arbitrage spreads with both trading volume and volatility. In the following section, we run regressions with spreads as the dependent variable and volatility and volume as independent variables. We show a strong positive relation between daily spreads and realised volatility.

To demonstrate that spreads and volatility have a strong positive relationship, we first calculate the daily-realised volatility. Realised volatility is the mean of the squared returns on a single day depicted by the equation below. As we use daily realised volatility, we also aggregate the spreads between each exchange to a daily level. We do this by taking the square of all of the spreads in a day and taking the sum. We use the square, as we are only interested in the absolute value of the spread, not whether it is positive or negative. Table 6 provides the regression results for each spread, and the equation below shows the regression. The exchange at the basis of the spread is also the one used to calculate the daily realised volatility, so we include all pairs to see the significance of using each exchange's price to calculate the volatility. As can be seen, over half of the regression had positive statistically significant betas to the 1% level, confirming that volatility has a positive impact on the spreads between the pairs. In some cases, the model produced an R^2 of 0.80, suggesting that volatility explains up to 80% in the variability of the spread. The model performed the best with pairs that included Bitstamp, suggesting that volatility had a more significant impact on Bitstamp's deviations in prices of Bitcoin than the other exchanges. Furthermore, those regressions that used the Bitcoin price from Bitstamp found that the realised volatility had higher explanatory power than those pairs that included only Bitstamp in the spread.

$$\hat{\sigma_t^2} = n^{-1} \sum_{i=1}^{n} r_{t-i}^2 \tag{11}$$

To investigate the causes of the large trading profits derived from pairs that include Coinbase, we consider the results of the regression. Regarding table 6 we see the majority of pairs including Coinbase had statistically insignificant relationships with volatility, except for those including Bitstamp. This result leaves us puzzled regarding the causes of the large profits. Comparing the results of the formation period with the trading period tells us that there is a change in the dynamics of Coinbase. In the formation period, pairs including Coinbase performed averagely whereas more recently they obtained the most profits. Therefore, the change in dynamics must be down to a new latent risk factor that predominately affects Coinbase. This latent risk factor would be a good place for succeeding research to concentrate.

$$Spr\hat{e}ad^2 = \beta_0 + \beta_1 \hat{\sigma_t^2} + \epsilon_i \tag{12}$$

Following this, we test the relationship between volume and spreads. We follow the same methodology as the previous regression. We take the total volume for a day and use the daily spread we calculated previously. Figure 14 shows both the daily volume and daily spread between

Gemini and Kraken. We multiplied the spread by 10,000 to aggregate it to a similar scale. The plot appears to show two variables that move closely together. However, the regression results tell a different story. The regression does not show a robust result between volume and spread, with at most volume explaining 13% in the variability of the spread. Furthermore, we do not find a single beta that is statistically significant. The result suggests that trading volume has not impacted on possible arbitrage profits. We believe this result may be because the exchanges included having minimal liquidity risk, and therefore the volume traded has minimal effect.

Figure 14: Relationship between Volume and Spread.

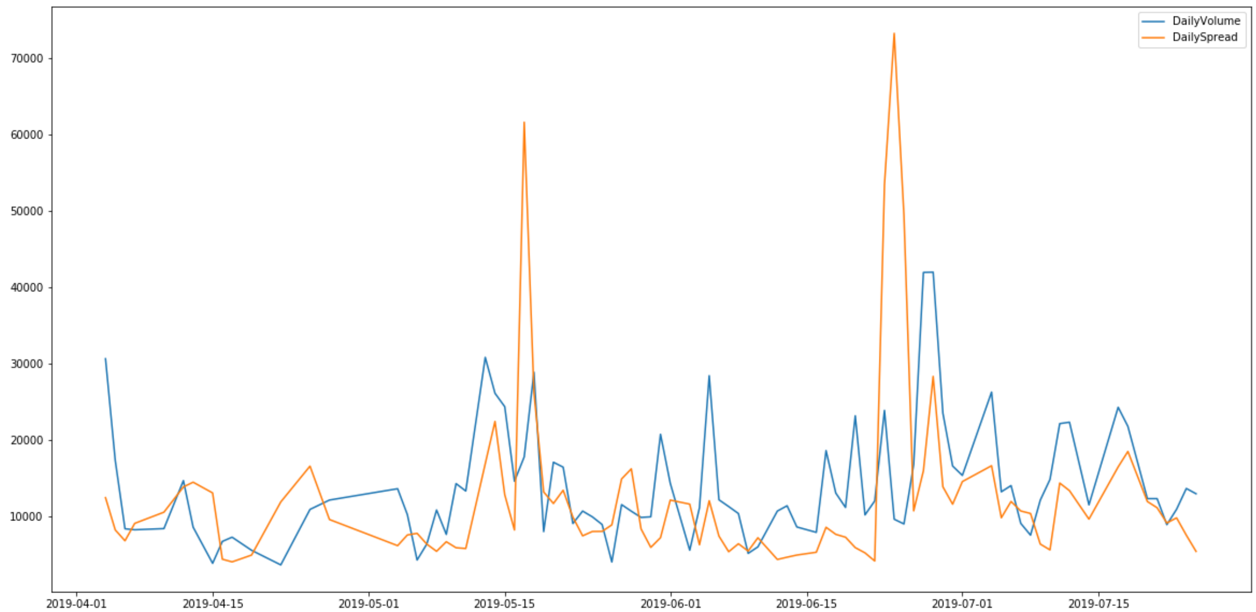

Figure 14 depicts the volume of Gemini and Kraken and the squared spread between the two exchanges. Note the spread is multiplied by 10,000

10 Conclusion

This research project studies arbitrage and market efficiency in cryptocurrency markets. We show that between 1 April 2019 and 14 August 2019 the dynamics of the deviations in prices of Bitcoin across exchanges change. We find that the start of the sample period exhibits low, stable spreads between exchanges; however, towards the end of the sample, we see much larger recurring price differences. This result contradicts a large majority of cryptocurrency literature, which suggests that Bitcoin markets are becoming more efficient.

Subsequently, we examine different trading strategies to see which best exploits the deviations across exchanges. We introduce a traditional pairs trading strategy which is most commonly used by hedge funds to exploit cointegrated stocks. We test both the traditional pairs trading strategy and a strategy put forward in cryptocurrency literature. For each strategy, we examine different changeable parameters to find what performs the best. We find that a traditional pairs trading strategy using an EqWMA Bollinger Band performs the best and provides a profit of $25,725.13 in our trading period. Furthermore, we find that those pairs, including Coinbase, provide the most profit within our trading period. This suggests that Coinbase diverges from the general Bitcoin price more than its counterparts do. We fail to find a reason for this; however, we suggest that it is not due to counterpart or liquidity risk. The reasons Urquhart (2016) suggest are cause discounts or premiums of exchanges. The resulting arbitrage profits contradict the work of Alexander and Heck (2019b) who find that arbitrage profits are eliminated once you include trading costs. We propose that the reopening of arbitrage profits is down to three factors; firstly, Bitcoin volatility increase towards the end of the sample following a stable period. Secondly, as Alexander and Heck (2019a) find CME futures have been used to manipulate Bitcoin spot prices; which may be a key factor in lessening market efficiency. Finally, due to a latent risk factor that predominantly affects Coinbase

References

Carol Alexander and Michael Dakos. A critical investigation of cryptocurrency data and analysis. *Available at SSRN 3382828*, 2019.

Carol Alexander and Daniel Heck. Price discovery, high-frequency trading and jumps in bitcoin markets. 2019a.

Carol Alexander and Daniel Heck. Price discovery and efficiency in bitcoin markets. 2019b.

Bitstamp-Lucas and Reddit.com. Bitstamp margin trading update. URL `https://www.reddit.com/r/Bitstamp/comments/841m7d/bitstamp_margin_trading_update/`.

Nicola Borri and Kirill Shakhnov. Cryptomarket discounts. *Available at SSRN 3124394*, 2018.

David A Bowen and Mark C Hutchinson. Pairs trading in the uk equity market: Risk and return. *The European Journal of Finance*, 22(14):1363–1387, 2016.

Binh Do and Robert Faff. Does simple pairs trading still work? *Financial Analysts Journal*, 66(4): 83–95, 2010.

Christopher Fink and Thomas Johann. Bitcoin markets. *Available at SSRN 2408396*, 2014.

Evan Gatev, William N Goetzmann, and K Geert Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 19(3):797–827, 2006.

Sashikanta Khuntia and JK Pattanayak. Adaptive market hypothesis and evolving predictability of bitcoin. *Economics Letters*, 167, 2018.

Andrew W Lo. Reconciling efficient markets with behavioral finance: the adaptive markets hypothesis. *Journal of investment consulting*, 7, 2005.

Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 2019.

Saralees Nadarajah and Jeffrey Chu. On the inefficiency of bitcoin. *Economics Letters*, 150, 2017.

Marcelo Scherer Perlin. Evaluation of pairs-trading strategy at the brazilian financial market. *Journal of Derivatives & Hedge Funds*, 15(2):122–136, 2009.

Andrew Urquhart. The inefficiency of bitcoin. *Economics Letters*, 148, 2016.

Andrew Urquhart. What causes the attention of bitcoin? *Economics Letters*, 166:40–44, 2018.

Faisal Nazir Zargar and Dilip Kumar. Informational inefficiency of bitcoin: A study based on high-frequency data. *Research in International Business and Finance*, 47:344–353, 2019.

Appendix

Table 1: *
The table provides the summary statistics of the Bitcoin index, based on 10-minute prices

Bitcoin Index	
Mean	8418.46
Standard Deviation	2430.03
Minimum	4076.96
Maximum	13849.8
Skew	-0.05906
Kurtosis	-1.29634

Table 2: *
The table provides the volume dependent fees as a percentage per transaction for all five exchanges.

Trading Volume	Coinbase	Bitstamp	Kraken	Gemini	ItBit
10,000	0.25	0.25-0.24	0.26-0.24	0.35	0.35
100,000	0.20	0.22-0.14	0.22-0.18	0.25	0.25
1,000,000	0.18	0.13-0.12	0.16-0.10	0.15	0.175

Table 3: *
The table provides the summary statistics of the Bitcoin price deviations across exchanges, based on 10-minute prices. The Average is denoted as in basis points and the St.Dev is annualized

	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Average	Coinbase	...	0.91	7.35	-1.29	-0.33
	Bitstamp	-0.95	...	-0.12	3.19	1.18
	Kraken	-8.12	0.41	...	-2.18	-1.19
	Gemini	1.26	2.18	2.18	...	0.95
	ItBit	0.31	1.16	1.19	-1.01	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
St. Dev.	Coinbase	...	0.47	0.45	0.41	0.27
	Bitstamp	0.47	...	0.25	0.25	0.30
	Kraken	0.45	0.24	...	0.20	0.25
	Gemini	0.41	0.25	0.20	...	0.24
	ItBit	0.27	0.30	0.25	0.24	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Min	Coinbase	...	-0.06	-0.07	-0.06	-0.03
	Bitstamp	-0.07	...	-0.09	-0.07	-0.07
	Kraken	-0.05	-0.01	...	-0.01	-0.03
	Gemini	-0.04	-0.09	-0.01	...	-0.03
	ItBit	-0.03	-0.03	-0.04	-0.03	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Max	Coinbase	...	0.07	0.05	0.04	0.03
	Bitstamp	0.06	...	0.01	0.01	0.03
	Kraken	0.06	0.08	...	0.01	0.03
	Gemini	0.06	0.07	0.09	...	0.03
	ItBit	0.30	0.07	0.03	0.030	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Skew	Coinbase	...	2.18	0.12	-1.40	-1.39
	Bitstamp	-3.24	...	-29.37	-19.93	-12.68
	Kraken	-0.91	25.05	...	0.06	-1.07
	Gemini	0.60	17.16	-0.09	...	-1.89
	ItBit	0.98	10.851	0.581	1.404	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Kurtosis	Coinbase	...	172.40	135.44	158.79	112.69
	Bitstamp	181.84	...	2331.54	1307.62	689.46
	Kraken	128.42	1873.41	...	8.09	147.23
	Gemini	41.35	1065.80	8.17	...	159.10
	ItBit	113.45	575.75	150.25	159.94	...

Table 4: *
The table provides the trading profits derived from each strategy considering all exchanges between 1 April 2019 and 7 June 2019 . The results are given in USD.

	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
ExWMA 2MSD	Coinbase	...	375.17	152.91	-0	86.97
	Bitstamp	375.17	...	239.91	181.20	254.38
	Kraken	152.92	238.91	...	84.58	108.09
	Gemini	0	181.20	84.58	...	79.17
	ItBit	86.97	254.38	108.09	79.17	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
ExWMA 1MSD.	Coinbase	...	204.513	84.63	0	86.97
	Bitstamp	204.513	...	65.63	206.47	28.90
	Kraken	84.63	65.63	...	84.58	132.79
	Gemini	0	206.47	84.58	...	53.85
	ItBit	86.97	28.90	132.79	53.85	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
EqWMA 2MSD	Coinbase	...	204.51	118.24	0	86.97
	Bitstamp	204.51	...	65.63	206.47	112.42
	Kraken	118.24	65.63	...	84.58	151.32
	Gemini	0	206.47	84.58	...	53.85
	ItBit	86.97	112.42	151.32	53.85	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
EqWNA 1MSD	Coinbase	...	34.89	84.63	0	181.31
	Bitstamp	34.89	...	65.63	39.23	28.90
	Kraken	84.63	65.63	...	51.54	260.11
	Gemini	0	39.23	51.54	...	184.46
	ItBit	181.31	28.90	260.11	184.46	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Traditional ExWMA	Coinbase	...	149.83	12.91	0	257.99
	Bitstamp	158.26	...	148.63	142.93	336.72
	Kraken	22.43	131.29	...	13.89	266.48
	Gemini	0	131.05	3.54	...	244.99
	ItBit	189.11	236.00	136.37	181.00	...
	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Traditional EqWMA	Coinbase	...	177.41	50.50	0	458.19
	Bitstamp	171.66	...	154.23	164.48	554.78
	Kraken	38.98	152.59	...	0	494.41
	Gemini	0	171.02	3.54	...	406.68
	ItBit	126.77	187.26	167.18	195.15	...

Table 5: *
The table provides the profits during our trading period from using a traditional pairs trading strategy and EqWMA Bollinger bands . The results are given in USD.

	Exchange	Coinbase	Bitstamp	Kraken	Gemini	ItBit
Trading Profits	Coinbase	...	4887.91	4432.36	2845.50	233.50
	Bitstamp	5218.72	...	111.29	37.11	71.07
	Kraken	4671.12	41.71	...	0	10.27
	Gemini	3005.36	8.73	0	...	-29.11
	ItBit	107.53	51.21	18.16	2.64	...

Table 6: *
Regression results with spread as the dependent variable and realized volume as the independent.

Variables:	intercept	β	p value	R^2
Coinbase Bitstamp	0.00029	20.27361**	0.00002	0.13204
		(4.63)		
Coinbase Kraken	0.00039	4.20941	0.25664	0.01045
		(3.69)		
Coinbase Gemini	0.00038	7.6435	0.15908	0.02266
		(5.38)		
Coinbase ItBit	0.00027	0.95822	0.77642	0.0026
		(3.34)		
Bitstamp Coinbase	0.0034	15.5270**	0	0.17309
		(3.02)		
Bitstamp Kraken	-0.0002	20.1289**	0	0.81491
		(0.85)		
Bitstamp Gemini	-0.00014	16.0103**	0	0.84431
		(0.73))		
Bitstamp ItBit	0	14.7831**	0	0.631243
		(2.02)		
Kraken Coinbase	0.0003	5.50593	0.19901	0.013376
		(4.26)		
Kraken Bitstamp	-0.0002	20.5196**	0	0.37113
		(2.38)		
Kraken Gemini	0	1.84849**	0.00084	0.12346
		(0.53)		
Kraken ItBit	0.00019	2.28164	0.54832	0.01213
		(3.75)		
Gemini Coinbase	0.00036	8.23063	0.13983	0.02488
		(5.52)		
Gemini Bitstamp	-0.0001	17.75657**	0	0.49961
		(1.89)		
Gemini Kraken	0.0009	1.63311*	0	0.11746
		(0.48)		
Gemini ItBit	0.0008	1.89644**	0	0.52852
		(0.34)		
ItBit Coinbase	0.00023	2.39722	0.48252	0.01603
		(3.37)		
ItBit Bitstamp	-0.0002	19.5086**	0.00002	0.44146
		(3.94)		
ItBit Kraken	0.00014	3.74387	0.23971	0.04576
		(3.12)		
ItBit Gemini	0.0008	1.93638**	0.0003	0.55872
			(0.33)	

**p<0.01, *p<0.05