Deviations from Triangular Arbitrage Parity in Foreign Exchange and Bitcoin Markets*

Julia Reynolds† Leopold Sögner‡ Martin Wagner§ Dominik Wied¶

This Version: November 5, 2018

Abstract

This paper applies new econometric tools to monitor and detect so-called "financial market dislocations," defined as periods in which substantial deviations from arbitrage parities take place. In particular, we focus on deviations from the triangular arbitrage parity for exchange rate triplets. Due to increasing media attention towards mispricing in the market for cryptocurrencies, we include the cryptocurrency Bitcoin in addition to fiat currencies. We do not find evidence for substantial deviations from the triangular arbitrage parity when only traditional fiat currencies are concerned, but document significant deviations from triangular arbitrage parities in the newer markets for Bitcoin. We confirm the implications of our results for portfolio strategies by showing that a currency portfolio that trades based on our detected break-points outperforms a simple buy-and-hold strategy.

Keywords: Triangular Arbitrage Parity, Foreign Exchange Markets, Cryptocurrencies, Cointegration, Monitoring

JEL Classification: G12, G15, C22, C32

*The authors appreciate helpful comments from Lin William Cong, Thomas Johann, Igor Pozdeev, Christian Wagner, Gloria Yu, and Jinyuan Zhang, and from conference and seminar participants at the University of Cologne, the 11th International Conference on Computational and Financial Econometrics (CFE 2017), the 2nd Vienna Workshop for High-Dimensional Time Series in Macroeconomics and Finance, the 30th Workshop of the Austrian Working Group on Banking & Finance (AWG), the 2017 ECDA Conference, the 2018 Bergen Fintech Conference, the XXVI Finance Forum, and the 2018 Meeting of the German Finance Association. We thank Etienne Theising for his help in collecting and compiling the Bitcoin data used in this analysis. Financial support from the Austrian Central Bank under Anniversary Grant Nr. 17053 is gratefully acknowledged. The authors gratefully acknowledge financial support by Deutsche Forschungsgemeinschaft via the Collaborative Research Center SFB823 (Projects A3 and A4). The views expressed do not necessarily reflect the views of the Bank of Slovenia/Banka Slovenije or the European System of Central Banks.

†Corresponding Author. Institute of Finance, Università della Svizzera italiana (Via Giuseppe Buffi 13, 6900 Lugano, Switzerland), `julia.elizabeth.reynolds@usi.ch`.

‡Department of Economics and Finance, Institute for Advanced Studies (Josefstädter Str. 39, 1080 Vienna, Austria), Vienna Graduate School of Finance and NYU Abu Dhabi, `soegner@ihs.ac.at`.

§Faculty of Statistics, TU Dortmund (Vogelpothsweg 87, 44227 Dortmund, Germany), Institute for Advanced Studies and Bank of Slovenia, `mwagner@statistik.tu-dortmund.de`.

¶Institute of Econometrics and Statistics, University of Cologne (Albertus-Magnus-Platz, 50923 Cologne, Germany), `dwied@uni-koeln.de`.

1 Introduction

Many empirical studies demonstrate that pricing identities following from no-arbitrage conditions – called *arbitrage parties* – often fail to hold exactly in real-world datasets. To explain such deviations, the limits to arbitrage literature considers restrictions or constraints on trading such that arbitrage opportunities cannot be exploited (see, e.g., Shleifer and Vishny, 1997; Gromb and Vayanos, 2010). Trading costs, illiquidity, and short-sale constraints are prominent examples of phenomena causing such market imperfections (see, e.g., Gagnon and Karolyi, 2010). In extreme cases, such as during periods of disaster or financial crisis, substantial arbitrage parity deviations can be observed for many asset types (see, e.g., Veronesi, 2004; Barro, 2006, 2009; Bollerslev and Todorov, 2011). Such periods of significant deviations from arbitrage parities are also referred to in the literature as "financial market dislocations," defined more precisely by Pasquariello (2014) as "circumstances in which financial markets, operating under stressful conditions, cease to price assets correctly on an absolute and relative basis."

Given the omnipresence of deviations from arbitrage parities even under normal market conditions, the empirical literature has turned to asking whether and under which circumstances these deviations are large enough to allow for riskless profits or represent market dislocations (see, e.g., Matvos and Seru, 2014; Fleckenstein et al., 2014; Pasquariello, 2014). The typical empirical method for calculating arbitrage parity deviations is to calculate them directly from prices and then check whether prevailing market conditions allow for arbitrage profits. However, this method requires the availability of quotation-level data, along with exact transaction costs, often for multiple markets – data that is notoriously difficult to obtain. Recent literature such as Akram et al. (2008) and Foucault et al. (2016) use high-frequency data to investigate deviations from the law of one price and the existence of arbitrage profits. Our methodology is based on the simple, and admittedly crude, idea that, under normal circumstances, no-arbitrage prices should share similar trends and observed arbitrage parity deviations should be stationary. Specifically, we consider deviations from the triangular arbitrage parity, which specifies a parity relationship between a triplet of currencies such that agents cannot profit from an instantaneous transaction between these currencies. Based on our observation that the triangular arbitrage parity does not hold exactly, we consider a stochastic version of (the log-form of) the triangular arbitrage parity by adding a noise term. This noise term captures non-zero deviations from the arbitrage parity, which are always observed. If the exchanges rates are integrated, the stochastic version of the parity corresponds to a cointegrating regression model, in which the parity implies specific values for regression

parameters. Our empirical methodology is sensitive to both changes in the time series properties of the noise terms (e.g., a transition from stationary to $I(1)$ or explosive) and to changes in the model parameters. That is to say, our analysis tests whether market forces result in significant deviations from a stochastic version of the triangular parity with stationary innovations and/or parameters implied by the exact parity. Our method for detecting financial market dislocations monitors structural breaks in time series properties, and is able to detect deviations even in aggregated, and low-frequency data.

To perform this analysis, we apply the monitoring tools of Wagner and Wied (2017). These tools use an expanding window detector to monitor residuals from a cointegrating regression for either a break from a cointegrating to a spurious regression, or a break in the parameters of the cointegration regression. This allows us to identify specific dates around which a break of either of the two forms in a cointegration regression takes place and, by examining historical circumstances around the dates of these structural breaks, to explore the variables and/or events that may have caused the financial market dislocation. In addition, we use bid-ask spreads to estimate no-arbitrage bounds, or limits within which the cost of an arbitrage trade is larger than its profit, to identify whether triangular arbitrage parity deviations can be explained by spread-implied transaction costs.

While deviations from the triangular arbitrage parity have been shown to occur in the market for fiat currencies (see, e.g., Lyons and Moore, 2009; Kozhan and Tham, 2012), they are usually rather rare. On the other hand, there is increasing media attention towards the potential for arbitrage in the market for *cryptocurrencies*.[1] We include Bitcoin, the most well-known and liquid cryptocurrency, in our analysis in addition to traditional fiat currencies. Arbitrage opportunities in the market for Bitcoin have also been examined by Dong and Dong (2014), who document persistent deviations of Bitcoin exchange rates from those implied by fiat currency markets, as well as Pieters and Vivanco (2017), who find evidence for substantial price deviations across Bitcoin exchanges. Makarov and Schoar (2018) show that there is substantial price dispersion across Bitcoin exchanges in different regions. Pieters (2016) and Yu and Zhang (2017) explore the relationship between triangular arbitrage parity deviations

[1] See, e.g., "Bitcoin's Crashing? That Won't Stop Arbitrage Traders From Raking in Millions," Wall Street Journal, 4 February 2018, available at `https://www.wsj.com/articles/bitcoins-crashing-that-wont-stop-arbitrage-traders-from-raking-in-millions-1517749201` and "As China cracks down, Japan is fast becoming the powerhouse of the Bitcoin market", CNBC News, 29 September 2017, available at `https://www.cnbc.com/2017/09/29/bitcoin-exchanges-officially-recognized-by-japan.html`. In addition, extreme cases of price disparities between cryptocurrencies in different markets can be observed. For example, on 23 October 2017, the Bitcoin exchange platform Golix reported prices of around \$9,600 per Bitcoin in the country of Zimbabwe, while at the same time Bitcoin traded at around \$6,000 per Bitcoin in U.S.-based Bitcoin exchanges (see "Bitcoin is breaking all kinds of price records in cash-strapped Zimbabwe", Quartz Africa, 23 October 2017, available at `https://qz.com/1109470/bitcoin-highest-price-is-in-zimbabwe-due-to-a-forex-shortage/`). Such disparities clearly violate the law of one price, potentially leaving room for arbitrage opportunities.

and cross-country differences in capital controls.

Our results show that break-points in deviations from the triangular arbitrage parity are rare when only fiat currencies are included in the currency triplet; we only find deviations in two (out of 36 possible) currency triplets involving the U.S. dollar and other fiat currencies, and in no cases against the euro. On the other hand, we detect a break-point in the majority of currency triplets that include Bitcoin. We show that our detected break-points correspond to major market events. Of particular importance to Bitcoin markets was the February 2014 bankruptcy and collapse of Mt. Gox, a Tokyo-based Bitcoin exchange that, at that time of its collapse, was handling around 70% of the world's Bitcoin trades (see, e.g., Decker and Wattenhofer, 2014). The bankruptcy resulted in the overnight loss of $473 million worth of Bitcoin, which substantially damaged investor confidence in Bitcoin (see, e.g., Fink and Johann, 2014). Our results show that particularly this event led to significant deviations from the triangular arbitrage parity for a wide range of currency triplets that include Bitcoin. Furthermore, while triangular arbitrage parity deviations for fiat currency triplets are virtually never outside of our measures of no-arbitrage bounds, deviations for Bitcoin currency triplets are often outside of no-arbitrage bounds implied by mean spreads. Overall, our results confirm the relative rarity of triangular arbitrage within the market for fiat currencies, while highlighting the potential for financial market dislocations on the newer market for cryptocurrencies. We confirm the potential implications of our results on portfolio strategies by showing that a currency portfolio that trades based on our detected break-points outperforms a simple buy-and-hold strategy.

Our dataset is described in Section 2, while Section 3 defines and provides more details on the triangular arbitrage parity. Section 4 describes the monitoring tools from Wagner and Wied (2017), used to detect deviations from the triangular arbitrage parity. Empirical results are presented in Section 5, and finally Section 6 concludes.

2 Data

Daily spot rates for fiat currencies are obtained from the Pacific Exchange Rate Service, which collects nominal noon spot exchange rates as observed and reported by the Bank of Canada.[2] These exchange rates represent averages of transaction prices and price quotes from financial institutions taken between 11:59am and 12:01pm Eastern time (ET). The sample includes exchange rates between the seven most

[2]The data is available through `http://fx.sauder.ubc.ca/`.

actively-traded currencies on foreign exchange markets (USD, EUR, JPY, GDP, AUD, CAD, CHF), as well as three relatively inactively-traded currencies (SEK, MXN, ZAR).

In addition, Bitcoin (XBT) exchange rates are obtained from *bitcoincharts*, a service that collects historical trade data from a cross-section of Bitcoin exchanges.[3] In order to match this data with our fiat currency sample, we collect all exchange-reported transaction prices and volumes between 11:59am and 12:01pm ET. Daily noon exchange rates are then calculated by taking the volume-weighted average transaction prices across exchanges. In this way, our Bitcoin rates should reflect the prices available to Bitcoin traders at relatively the same time as our fiat currency observations. We collect Bitcoin exchange rates against all fiat currencies in our sample for which data is available: USD, EUR, AUD, CAD, CHF, GBP, JPY, SEK, ZAR, MXN. To limit the effects of extreme prices, a minimum of three transaction prices is required to calculate the noon exchange rate. More details about the fiat and Bitcoin exchange rate data, including their time series properties, are provided in Appendix A.1.

The sparser nature of Bitcoin trading leaves us with missing values. To treat missing values, we proceed as follows: If less than three transactions are found to occur between 11:59am and 12:01pm ET, our algorithm then takes the next-closest observations to the target window in terms of time, until a minimum of three observations are found. More details regarding this procedure and the resulting precision of Bitcoin rates can be found in Appendix A.2.

Figure 1 plots exchange rates relative to the USD during our sample time period, which lasts from 1 May 2013 until 31 December 2015.[4] Note that the XBT/USD exchange rate exhibits much higher volatility than the traditional fiat currencies. The largest peak corresponds to 29 November 2013, after a surge in retailer announcements that they would soon be accepting Bitcoin as payment,[5] and after a hearing by the U.S. Senate demonstrated that they would take a more neutral position regarding digital currencies.[6]

[Insert Figure 1 about here.]

[3]The complete trade history data is available publicly at `http://api.bitcoincharts.com/v1/csv/`. In order to be listed on bitcoincharts, exchanges voluntary submit their complete trade and orderbook history to bitcoincharts. Similar datasets are used in Fink and Johann (2014), Dong and Dong (2014) and Pieters and Vivanco (2017). Note that Bitcoin exchanges are typically operating 24 hours a day.

[4]The sample time period is chosen to allow for sufficient liquidity in Bitcoin markets.

[5]See, e.g., "Online retailer Overstock.com plans to accept Bitcoin payments next year", yahoo! Finance, 21 December 2013, available at `https://finance.yahoo.com/news/online-retailer-overstock-com-plans-accept-bitcoin-payments-183244697.html`.

[6]See, e.g., "Why Regulate Bitcoin?", The New Yorker, 18 November 2013, available at `https://www.newyorker.com/business/currency/why-regulate-bitcoin`.

To calculate our measures of no-arbitrage bounds, we collect fiat currency bid and ask quotes provided by the exchange rate service WM/Reuters for the time period 1 May 2013 until 31 December 2015. Bid and ask quotes are available for 60 out of 90 possible combinations of fiat currencies. The quotes are downloaded from Thomson Reuters Datastream, and are originally sourced from wholesale electronic currency platforms, including Thomson Reuters Matching, EBS, and Currenex, to reflect an average of interbank quotes between 3:59:30pm and 4:00:30pm GMT (10:59:30am and 11:00:30am ET).[7] Defining the absolute spread by $\kappa_{t,A/B} := S^a_{t,A/B} - S^b_{t,A/B}$, the percentage bid-ask spread is calculated as $\delta_{t,A/B} := \frac{\kappa_{t,A/B}}{S^a_{t,A/B}} \times 100$.

Additionally, daily percentage bid-ask spreads for Bitcoin for the same time period are downloaded from *Bitcoinity*; the data includes percentage bid-ask spreads for all exchange rates except for against the South African rand and Mexican peso.[8] Bitcoinity calculates percentage bid-ask spreads as the percentage difference between the daily *minimum* ask quote and *maximum* bid quote; thus, these spreads do not represent a spread that would be available to a trader at a particular time, but should represent an upper bound on bid-ask spreads in the Bitcoin market.

Table 1 presents summary statistics for percentage bid-ask spreads for the nine fiat exchange rates against the USD, as well as the eight exchange rates against Bitcoin. From the table, it is clear that Bitcoin spreads tend to be much higher than those of fiat currencies. For example, the percentage spread for trading EUR against XBT is about 62 times larger than the spread for trading EUR against USD. Figure 2 plots the average percentage bid-ask spreads for exchange rates involving fiat currencies only and the average percentage bid-ask spreads for Bitcoin exchange rates over our sample time period. The plot again highlights the difference in magnitude between the two groups of spreads, and also shows that they tend to follow different dynamics. Spreads for Bitcoin begin to rise in early February 2014, i.e., right around the Mt. Gox bankruptcy, spiking at 12.4% several weeks later on 14 March 2014. Interestingly, Bitcoin spreads peak again on 5 November 2015, following a surge in Bitcoin's value.[9]

[Insert Table 1 and Figure 2 about here.]

[7] On December 14, 2014, WM/Reuters moved to using a five-minute, rather than a 60-second, window, around 4pm GMT. See `https://www.reuters.com/article/markets-forex-fixings/london-forex-fix-moves-to-5-minute-window-on-dec-14-wm-memo-idUSL6N0T73BW20141117`.

[8] The data is available from `data.bitcoinity.org/`. Bid-ask spreads are provided separately for each Bitcoin exchange. We take the maximum bid-ask spreads across exchanges.

[9] See, e.g., `https://www.coindesk.com/bitcoin-hits-450-for-the-first-in-2015/`.

3 The Triangular Parity

In modern financial theory, one consequence of an arbitrage-free market is that the prices of related assets should be fundamentally linked, such that we should see arbitrage parities between the prices of related assets.[10] In particular, consider three currencies A and B, as well as V, where V denotes the so-called "vehicle currency." Let $S_{t,A/B}$ denote units of currency A received for one unit of currency B in period t. Consider a capital market that is absent trading constraints and costs, in which agents are perfectly informed and have no market power (*perfect capital market*). In the absence of arbitrage, for any triplet of spot exchange rates $S_{t,A/B}$, $S_{t,A/V}$ and $S_{t,V/B}$, we obtain the *triangular arbitrage parity*:

$$S_{t,A/B} \quad = \quad S_{t,A/V} S_{t,V/B} \text{ or, equivalently, } \ln S_{t,A/B} = \ln S_{t,A/V} + \ln S_{t,V/B} \ . \tag{1}$$

In a hypothetical frictionless and arbitrage-free market, we should always observe $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/B} = 0$ exactly. However, for our empirical exchange rate data described in Section 2, we observe $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/B} \neq 0$ for all periods t and all currency triplets in our sample.

Under "normal" market conditions, small deviations from arbitrage parities are expected to persist, for example, due to transaction costs. The existence of transaction costs create so-called "no-arbitrage bounds," within which the cost of an arbitrage trade is larger than its profit and thus these deviations persist in the market (see, e.g., Modest and Sundaresan, 1983; Klemkosky and Lee, 1991; Engel and Rogers, 1996). Therefore, these deviations do not necessarily constitute tradable arbitrage opportunities in practice. Whether our observations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/B} \neq 0$ can be explained by transactions costs implied by bid-ask spreads will be explored later in this section.

Figure 3 presents deviations from the triangular arbitrage parity for a representative sample of currency triplets including only fiat currencies, in which the U.S. dollar (USD) is used as the vehicle currency.[11] These deviations are quite small in magnitude, remaining in the interval $[-0.0001, 0.0001]$ or even smaller, and are fluctuating around zero. This holds even for currencies such as the MXN and ZAR, which likely have higher transactions costs due to their relatively low trading volumes. For comparison, Figure 4 presents all deviations for triplets including Bitcoin, in which USD is used as the vehicle currency. It is clear that including Bitcoin into the currency triplets leads to both higher

[10] Following Skiadas (2009), Definition 5.1, consider an adapted process of payoffs $\{c(\omega, t)\}$, $t = 0, 1, \ldots, T$ and $\omega \in \Omega$, also called cash flow. A cash flow where $c(\omega, t) \geq 0$ for all ω, t and $c(\omega, t) > 0$ for at least one ω, t is called an *arbitrage*. A market that does not admit arbitrage is said to be *arbitrage-free*.

[11] Deviations for fiat currency triplets in which the euro is used as the vehicle currency look very similar.

magnitudes of and higher volatility in the deviations, as some are shown to leave the interval $[-1, 1]$. Furthermore, we observe time spans in which the deviations remain systematically smaller or larger than zero, pointing to persistent over- and underpricing in the market for Bitcoin. Such mispricings may persist, for example, due to the relatively low rate of informed institutional trading on Bitcoin exchanges.[12] Interestingly, most Bitcoin triplets experience a spike in their deviations around the Mt. Gox bankruptcy on 24 February 2014 (marked by the blue dotted line in the plots), showing the impact that this event had on the market for Bitcoin. By contrast, for fiat currency triplets we do not observe spikes around this date.

[Insert Figures 3 and 4 about here.]

We then investigate whether the deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ stay within no-arbitrage bounds implied by transaction costs as measured using bid-ask spreads. As the difference between bid and ask prices represents the profit earned by brokers to facilitate the market between buyers and sellers, the bid-ask spread represent a significant hidden cost for a round-trip trade, and therefore the bid-ask spread is often used as a measure of transaction costs. To calculate no-arbitrage bounds, we calculate upper and lower bounds for the bid-ask spread, as well as approximate measures of the mean location of the spread. In particular, as described in Appendix A.3, we calculate non-arbitrage bounds based on a measure of the maximum spread, $\bar{s}$, a 90% percentile measure, $s_{90\%}$,[13] mean and median spread measures, s_{mean} and s_{median}, and a measure for the minimum spread, $\underline{s}$.

Table 2 presents relative frequencies at which the deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ exceed our bid-ask spread-implied measures of no-arbitrage bounds. For all fiat currency triplets, we find that the deviations are never larger than our estimates of transaction costs, meaning that these deviations stay well within the no-arbitrage bounds implied by bid-ask spreads.[14] On the other hand, Bitcoin

[12] Anecdotal evidence puts the amount of Bitcoin held by institutional investors at about 1% as of 2017 (see, e.g., `https://bravenewcoin.com/news/bitcoin-2018-show-me-the-institutional-money/`). Furthermore, the most recent data from Bitcoin trading services provider itBit shows that institutional investors make up a majority of OTC Bitcoin transactions, implying that when they do trade, they may prefer to do so off-exchange (see `https://www.itbit.com/blog/itbit-bitcoin-otc-market-recap-may-2016-infographic`).

[13] Since the Bitcoin bid-ask spreads are obtained from various exchanges and include some extreme realizations, we consider the 90% percentile of the bid-ask spread distribution to exclude possible effects arising from such spikes in the bid-ask spread.

[14] One caveat to this result might be the different time windows use for the calculation of spot rates (12pm ET) and bid and ask quotes (4pm GMT/11am ET). Studies such as Marsh et al. (2017) and Evans (2018) have focused on the fact that fiat currency markets display vastly different market microstructure dynamics during the "4pm London Fix", as large banks flood the market with liquidity in attempts to manipulate this benchmark spot rate. However, these studies show that spreads during the 4pm GMT window are much *lower* than at other times during the trading day. This would therefore bias our results towards finding more deviations outside the bounds.

currency triplets are shown to exceed the no-arbitrage bounds much more often. This is particularly striking, considering that the maximum spread measure is much higher for Bitcoin (see Table 1). Deviations in excess of the maximum bid-ask spread remain relatively rare (less than 1%), with the exception of the triplet JPY-XBT-USD. However, a substantial proportion of deviations are above the upper bounds implied by the 90% percentiles of bid-ask spreads. Even for the most liquid Bitcoin triplet (EUR-XBT-USD), 22% of deviations are shown to exceed $\bar{s}_{90\%}$.

Therefore, it may be that arbitrage opportunities exist in the market for Bitcoin. However, it may also be that there are additional trading costs or frictions for Bitcoin that are not captured by bid-ask spreads, which limit the profitability of correcting deviations from the triangular arbitrage parity. Such trading costs may come, for example, from high exchange fees; Kim (2017) estimates a maximum Bitcoin exchange fee of about 0.5%, while Fink and Johann (2014) document Bitcoin exchange fees up to 2%. Another factor is the high latency of bitcoin trading; Courtois et al. (2014) show that Bitcoin transactions can take up to ten minutes, making it difficult to engage in simultaneous transactions with lower latency fiat currency markets. Bitocin traders can opt for more speed, but this requites the payment of a so-called "mining fee," which Fink and Johann (2014) estimate as about 1%-4% for Bitcoin. There are also a number of additional risks that are relatively unique to Bitcoin – such as higher risks of exchange insolvency and theft (see, e.g., Moore and Christin, 2013), or its association with illegal activity (see, e.g., Foley et al., 2018).

[Insert Table 2 about here.]

After examining deviations from the log-form of the triangular arbitrage parity (i.e., $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/B}$) and a brief comparison of these deviations to the no-arbitrage bounds implied by bid-ask spreads, the next section will analyze whether these deviations are connected to significant market dislocations, in terms of deviations from a stochastic version of the triangular arbitrage parity with stationary innovations and parameter values as implied by the exact parity.

4 Monitoring Market Dislocations

We consider the parity described in (1) as a cointegrating regression:

$$\begin{aligned} \ln S_{t,A/B} &= \alpha + \boldsymbol{\beta}^\top \begin{pmatrix} \ln S_{t,A/V} \\ \ln S_{t,V/B} \end{pmatrix} + u_t, \\ y_t &= \alpha + \boldsymbol{\beta}^\top \mathbf{x}_t + u_t \ , \end{aligned} \tag{2}$$

in which the spot rates $S_{t,A/B}$, $S_{t,A/V}$ and $S_{t,V/B}$ are assumed to be integrated. Again, in a hypothetical frictionless and arbitrage-free market, $\alpha = 0 =: \alpha^*$, $\boldsymbol{\beta} = (1,\ 1)^\top := \boldsymbol{\beta}^*$ and $u_t = 0$ for all $t = 1, \ldots, T$.

The following analysis applies *stationarity monitoring*, i.e., monitors the residuals $\widehat{u}_t$ for a change in cointegration behavior.[15] Let $[mT]$ denote the integer part of mT, where $m \in (0,1)$. Set $\boldsymbol{\beta} = \boldsymbol{\beta}^* = (1,1)^\top$ and consider $u_t = y_t - \alpha - \boldsymbol{\beta}^{*\top}\mathbf{x}_t$ as well as the *ordinary least squares* (OLS) residuals $\widehat{u}_t := y_t - \widehat{\alpha}_m - \boldsymbol{\beta}^{*\top}\mathbf{x}_t$. We denote by $\widehat{\alpha}_m$ the OLS estimate of α based on the regression model:

$$\left(y_t - \boldsymbol{\beta}^{*\top}\mathbf{x}_t\right) = \alpha + u_t \ . \tag{3}$$

The sub-sample $(y_t, \mathbf{x}_t :\ t = 1, \ldots, [mT])$ is used as a calibration period in order to estimate α as well as the long-run variance $\omega^2 := \sum_{j=-\infty}^{\infty} \mathbb{E}\left(u_{t-j}u_t\right)$.

To test the null hypothesis of "no structural change", i.e., that u_t is integrated of order zero ($I(0)$) and $\boldsymbol{\beta} = (1,\ 1)^\top$ for all $t = 1, ..., T$, Wagner and Wied (2017) consider an expanding window detector based on the statistic

$$\widehat{H}^m(s) := \frac{1}{T^2}\frac{1}{\widehat{\omega}_m^2}\sum_{j=[mT]}^{[sT]}\sum_{t=1}^{j}\widehat{u}_t^2 \ , \tag{4}$$

where $s \in [m,1] \subset (0,1]$ and $\widehat{\omega}_m^2$ is an estimate of the long-run variance matrix ω^2, estimated from the calibration period $t = 1, ..., [mT]$.[16] Under the assumptions listed in Wagner and Wied (2017), it holds

[15] An alternative to stationarity monitoring is cointegration monitoring, in which case the parameters $\boldsymbol{\beta}$ are also estimated using modified least squares methods. For more details, see Appendix A.4 and Wagner and Wied (2017).

[16] Estimates of the long-run variance matrix ω^2 are obtained by applying the Bartlett kernel. The bandwidth chosen follows from Andrews (1991)[p. 834].

under the null hypothesis for $T \to \infty$ that:

$$\widehat{H}^m(s) \Rightarrow H^m(s) := \frac{1}{\omega^2} \int_m^s \omega^2 \widetilde{W}(r)^2 ds = \int_m^s \widetilde{W}(r)^2 ds \ , \tag{5}$$

where $\widetilde{W}(s)$ is a demeaned Brownian motion and $\Rightarrow$ denotes weak convergence.

The test statistic $\widehat{H}_m(s)$ can be obtained for each $s \in (m, 1]$, such that $[mT] + 1 \leq [sT] \leq T$. A break-point τ_m is detected as the first s at which the test statistic $\widehat{H}^m(s)$ exceeds a critical value $c(\alpha_c, g)$. In particular, given a significance level α_c, Wagner and Wied (2017) define

$$\tau_m := \left\{ \min_{s:\ [mT]+1 \leq [sT] \leq T} \left| \frac{\widehat{H}^m(s)}{g(s)} \right| > c(\alpha_c, g) \right\} \ , \tag{6}$$

where $g(s)$ is a continuous weighting function. We follow Wagner and Wied (2017) and set $g(s) = s^3$ in the intercept-only case. The critical values $c(\alpha_c, g)$ follow from $\mathbb{P}\left(\sup \left| \frac{H^m(s)}{g(s)} \right| > c\right) = \alpha_c$ and can be obtained by means of simulations. A graphical illustration of the monitoring procedure is provided in Figure 5, where a break is detected in period $t_{\tau_m} := [\tau_m T]$.

[Insert Figure 5 about here.]

5 Results

This section describes the results from our empirical investigation of the triangular arbitrage parity in the foreign exchange and Bitcoin markets. Recall that the triangular arbitrage parity examines the relationship between a triplet of exchange rates. As in Pasquariello (2014), we include only permutations in which the USD or EUR are used as the vehicle currency. Furthermore, since $S_{t,A/B} \approx 1/S_{t,B/A}$, only one permutation of the same three currencies is included in the sample. This leaves us with $720/10 \div 2 = 36$ possible permutations of fiat currency triplets, and nine currency triplets involving Bitcoin, for each of the two vehicle currencies, thus giving us a total of 72+18=90 currency triplets.

Recall that our monitoring procedure detects instabilities in the noise u_t, and also in the parameters α and $\boldsymbol{\beta}$ from the regression in (2). To get a first look at the stability of the parameter estimates over time, we first use rolling window regressions to visually inspect whether the parameter estimates are close to the values implied by the triangular arbitrage parity, i.e., $\alpha = 0$ and $\beta_1 = \beta_2 = 1$. We consider rolling window parameter estimates for the regression model (2), using a window size $m = 0.2$

as a proportion of the total time series length. We estimate the parameters using fully modified least squares (FM-OLS), which corrects for possible serial correlation as well as endogeneity in a cointegration regression (see Phillips and Hansen, 1990), and abbreviate the estimated parameters by $\widehat{\alpha}^{+}_{mR}$ and $\widehat{\boldsymbol{\beta}}^{+}_{mR}$.

Representative results for the estimates $\widehat{\alpha}^{+}_{mR}$ and $\widehat{\boldsymbol{\beta}}^{+}_{mR} = \left(\widehat{\beta}^{+}_{1,mR}, \widehat{\beta}^{+}_{2,mR}\right)^{\top}$ are provided in Figure 6. The first two rows of these figures show rolling window estimates for currency triplets in which only fiat currencies are included, showing that the estimated parameters hardly deviate from their theory-implied values over our sample period, and that the variation of the estimates is small. On the other hand, the rolling window estimates for currency triplets involving Bitcoin, shown in the third and fourth rows of Figure 6, show high variation in the estimates and large deviations away from the theory-implied parameter values.[17] This is particularly the case for the estimates of α and β_1.[18]

[Insert Figure 6 about here.]

In a next step, we perform stationarity monitoring as described in Section 4, which will also allow us to identify dates around which structural breaks-points are more likely to have occurred. Overall, given the frequent observation of deviations outside of no-arbitrage bounds for currency triplets involving Bitcoin (see Figure 4), as well as deviations of parameter estimates away from their theory-implied values from the rolling window regressions (see Figure 6), we might expect that our monitoring tool will be more likely to reject the null hypothesis of "no structural breaks" in the cointegrating relationship for Bitcoin currency triplets.

In our stationary monitoring procedure, we fix $\boldsymbol{\beta} = \boldsymbol{\beta}^* = (1,1)^{\top}$ and monitor the OLS residuals $\widehat{u}_t := y_t - \widehat{\alpha}_m - \boldsymbol{\beta}^{*\top}\mathbf{x}_t$, with $m = 0.2$. As our sample period ranges from 1 May 2013 until 31 December 2015, the choice of $m = 0.2$ means that our calibration time period lasts from 1 May 2013 until 8 November 2013, which corresponds to a relatively stable period for exchange rates. The detected break-points are presented in Table 3. The results indicate that we are much more likely to detect a break-point in the triplets that include Bitcoin. With the USD as the vehicle currency, we observe six break-points in the triplets that include Bitcoin (XBT) (i.e., in two-thirds of the Bitcoin triplets), and only two break-points when the triplet is composed of fiat currencies only (5.5% of the fiat-only

[17] The results from some Wald-type tests based on the calibration sample and under the null hypothesis that the parameters are equal to the values implied by the arbitrage parity are provided in Appendix A.4. Similarly, these results show that the null is rejected much more often for currency triplets that include Bitcoin, as compared to currency triplets that only involve fiat currencies.

[18] Similar results are obtained when including a linear trend.

triplets).[19] Similarly, using the EUR as the vehicle currency leads to seven break-points in the Bitcoin triplets, and to zero breaks in the fiat-only triplets.

[Insert Table 3 about here.]

As we have shown that a high number of break-points are detected in the Bitcoin triplets, the next question is whether the detected break-points tend to be clustered in time. Figures 7 and 8 plot the Bitcoin exchange rates, with red vertical lines representing the detected break-points for currency triplets using the USD and EUR, respectively, as vehicle currencies. The detected break-point of 26 February 2014 for the triplet JPY-XBT-USD comes just days after the Mt. Gox bankruptcy. That the bankruptcy should hit the market for JPY/XBT the hardest is no surprise; the shut-down of the Tokyo-based exchange represented a massive disruption to the market for Bitcoin against the yen, as Bitcoin dropped 17.41% against the yen within a single day. Furthermore, as Mt. Gox is the only exchange in our dataset that offered JPY/XBT transactions at the time, its bankruptcy meant a virtual halt on exchange-based trading of yen for Bitcoin.[20]

However, other break-points for USD triplets do not seem to be clustered in time, although all take place after the Mt. Gox bankruptcy. It is not clear whether these later detection dates are due to a lagged reaction of our detector, or reactions to other market events. For example, it is likely that the the break-point of 19 January 2015 for the triplet CHF-XBT-USD is a result of the Swiss National Bank's abrupt ending to its cap on the CHF/EUR exchange rate just a few days earlier.[21]

For the triplets using EUR as a vehicle currency, depicted in Figure 8, interestingly we see that the break-points indeed tend to be clustered in time, with most break-points clustered around the bankruptcy of Mt. Gox and the subsequent recovery. The only exception it JPY-XBT-EUR; for this triplet, we already detect a break-point on 29 January 2014, *before* the Mt. Gox bankruptcy. Mt. Gox was experiencing problems throughout early 2014 in the lead-up to its bankruptcy, as customers complained about withdrawal delays and poor service.[22] Meanwhile, the break-points for almost all

[19] Interestingly, the two breaks in the fiat-only triplets both include the Mexican peso (MXN), with break-point dates in October 2014. The Mexican peso experienced a rapid decline at the end of 2014 due to a drop in oil prices and strengthening of the U.S. dollar, which might have generated these deviations. See, e.g., "Mexican Peso Falls to Two-Year Low as Central Bank Dims Outlook", Bloomberg, 6 December 2014, available at `https://www.bloomberg.com/news/articles/2014-12-05/mexican-peso-falls-to-five-year-low-as-central-bank-dims-outlook`.

[20] Note that there may have been more exchanges offering JPY/XBT that have chosen not to self-report in the bitcoincharts dataset, and furthermore that traders could search for counterparties in JPY/XBT via OTC markets.

[21] See, e.g., "Swiss Franc Soars After Central Bank Drops Cap", The New York Times, 15 January 2015, available at `https://www.nytimes.com/2015/01/16/business/swiss-national-bank-euro-franc-exchange-rate.html`.

[22] See, e.g., `https://www.coindesk.com/poll-mt-gox-withdrawal-issues/`.

remaining Bitcoin triplets occur during the weeks and months after the Mt. Gox bankruptcy, when trading was likely much more restricted due to the collapse of the major exchange. Our results imply that these trading restrictions may have led to deviations from the triangular arbitrage parity for Bitcoin traders in a variety of currencies. However, given difficulties in trading, it is possible that Bitcoin traders were not able to take advantage of these potential arbitrage opportunities.

[Insert Figures 7 and 8 about here.]

Figures 9 and 10 plot the OLS residuals $\widehat{u}_t := y_t - \widehat{\alpha}_m^\top - \boldsymbol{\beta}^{*\top}\mathbf{x}_t$ as described in Section 4, for two examples of triplets: one in which a break-point is detected (JPY-XBT-USD), and one in which no break-point is detected (JPY-SEK-USD). In the figures, the black dotted lines correspond to the cut-off between the calibration and monitoring time periods, and the red dotted lines correspond to the detected break-point (if any) for the currency triplet shown. In Figure 9, the break-point of 26 February 2014 corresponds to a downward spike in the residuals; furthermore, after the detected break-point, the residuals remain consistently negative. For comparison, Figure 10 shows a triplet for which no break-point is found. These residuals remain clustered around zero.

[Insert Figures 9 and 10 about here.]

Note that a break-point detected by means of our monitoring tool does not automatically imply that exploitable arbitrage opportunities exist. The deviations from triangular arbitrage could be due to costs and trading constraints that would prevent a profitable trade based on the deviations. Small deviations in α from the theoretical value $\alpha^* = 0$ can be justified by rounding errors or measurement effects. However, a break-point detected by our monitoring tool implies a break down in the cointegrating relationship between exchange rates, and thus indicates a substantial dislocation between the markets for these currencies.

5.1 Portfolio Trading Strategy

In a next step, we check whether our results could be used to construct a profitable trading strategy. The idea is to compare two portfolios: one based on a simple buy-and-hold strategy, to a portfolio that trades once a break in triangular arbitrage parity deviations is detected according to the Wagner and Wied (2017) monitoring procedure. If detected break-points indeed represent market dislocations, then a break may signal higher expected uncertainty in currency markets. Thus, it may be profitable to

re-balance the portfolio away from the implicated currencies. Similar strategies are explored in Wied et al. (2012, 2013). This strategy is furthermore similar to the portfolio re-balancing channel of Hau and Rey (2006), in which investors shift their portfolios away from foreign currencies when their exposure to exchange rate risk increases.

We compare two very simple portfolios that initially divide 100 units of the domestic currency DOM equally between the eleven different currencies in our sample: USD, EUR, JPY, GDP, AUD, CAD, CHF, SEK, MXN, ZAR, and XBT. In the following analysis we consider either the U.S. dollar, euro, Swiss franc, or Japanese yen as the domestic currency, $DOM \in \{USD, EUR, CHF, JPY\}$. These currencies are chosen as they are the so-called "hard" or "safe haven" currencies most often held by international investors (see, e.g., Campbell et al., 2003, 2010).[23] Our comparison excludes trading cost and assumes that any trade can be executed at the average prices described in Section 2. The initial and liquidation dates for both portfolios are defined as the first and last dates of the monitoring period in our sample, such that the durations of the portfolios are from 11 November 2013 until 31 December 2015, i.e., $t = [mT] + 1, ..., T$.

The first portfolio (denoted by Portfolio A) consists of a simple buy-and-hold strategy that never trades the currencies after the initial purchase. In this portfolio, the number of units $s^A_{t,i}$ of currency i held in the portfolio are held fixed after the initial purchase, such that $s^A_{t,i} = \frac{100}{11} S_{[mT]+1,i/DOM}$ for all $i = 1, \ldots, 11$ and $t = [mT] + 1, \ldots, T$. For example, if the domestic currency DOM is USD, and initial spot rate of JPY for USD is $S_{[mT]+1,JPY/USD} = 99.6393$, then the portfolio will initially purchase $s^A_{[mT]+1,JPY} = (100/11) * S_{[mT]+1,JPY/USD} = 905.8119$ units of Japanese yen. Note that the portfolio will hold $s^A_{[mT]+1,USD} = 100/11$ units of the domestic currency, which will obviously have a constant exchange rate of one. We assume that currencies are deposited and earn the local risk-free rate (assumed to be the one-month deposit rate, as collected from Thomson Reuters Datastream and transformed to a daily interest rate). As there is no risk-free Bitcoin-denominated asset, we assume that the risk-free rate for Bitcoin is zero. The buy-and-hold return for this portfolio assumes that the units of foreign currency are exchanged for the domestic currency on the liquidation date, and is thus defined as:

$$\mathcal{R}^A_{[mT]+1:T} := \frac{\left(\sum_{i=1}^{11} s^A_{T,i} S_{T,DOM/i} - 100\right)}{100} , \tag{7}$$

[23] An additional approach could consider investments into foreign assets, rather than foreign currencies. We opted to consider foreign currency investments because, first, it isolates the effect of exchange rate risk, i.e., the introduction of equity risk is avoided; and secondly, there are currently no widely accessible Bitcoin-denominated assets.

where $s^A_{T,i}$ is the terminal value of the initial units purchased of currency i at time T, $i = 1, \ldots, n$, defined as $s^A_{T,i} = s^A_{[mT]+1,i} \cdot \prod_{w=0}^{T-1}(1+\mathcal{R}^i_{f,w:w+1})$, where $\mathcal{R}^i_{f,t:t+1}$ is the daily risk-free rate for currency i. $S_{T,i/DOM}$ are the spot exchange rates of the currencies $i = 1, \ldots, n$ for DOM on the liquidation date T. Note that changes in the value of the buy-and-hold portfolio A are, in addition to compound interest, caused by currency revaluations within the time span $[mT] + 1$ to T. Hence, $\mathcal{R}^A_{[mT]+1:T}$ measures the mean performance of a domestic currency against the other ten currencies considered, after accounting for compound interest.

The second portfolio (denoted by Portfolio B) trades on the day(s) that the Wagner and Wied (2017) monitoring procedure detects a break-point in the triangular arbitrage parity deviations; specifically, it exchanges the foreign currencies involved in the triangular arbitrage triplet implicated by the break-point for the domestic one. Initially, Portfolio B purchases the same number of initial units as Portfolio A, such that $s^B_{t,i} = s^A_{t,i} = \frac{100}{11} S_{[mT]+1,i/USD}$ for all $i = 1, \ldots, 11$ at $t = [mT]+1$. Let $\tau_{m,k}$ denote the kth break-point detected for some currency triplet. The set $I_{\tau_{m,k}}$ collects those currency indices involved in the kth break but not involved in the former breaks $\tau_{m,1}, \ldots, \tau_{m,k-1}$. The investor will hold $s^B_{t,j} = 0$ of the foreign currencies j involved in the break-point $\tau_{m,k}$ (i.e. $j \in I_{\tau_{m,k}}$ or $j \in I_{\tau_{m,\kappa}}$, $\kappa < k$), while $s^B_{t,j} = s^A_{t,j} = s^A_{[mT]+1,j}$ for those foreign currencies j that have not yet been involved in any break. If currency j has already been involved in a break, $s^B_{t,j} = 0$ at the first t this currency was involved in one of the breaks $\tau_{m,1}, \ldots, \tau_{m,k-1}$. If at least one break is detected, that is $k \geq 1$, then at any t, fulfilling $\tau_{m,k} \leq t \leq \min\{T, \tau_{m,k+1} - 1\}$, the number of units of the domestic currency DOM will be equal to $s^B_{t,DOM} = s^A_{[mT]+1,DOM} \cdot \prod_{w=0}^{t-1}(1 + \mathcal{R}^{DOM}_{f,w:w+1}) + \sum_{\kappa=1}^{k} \sum_{j \in I_{\tau_{\kappa m}}} s^A_{[mT]+1,j} \cdot S_{\tau_{m,\kappa},DOM/j} \cdot \prod_{w=0}^{\tau_{m,k}-1}(1 + \mathcal{R}^j_{f,w:w+1}) \cdot \prod_{w=\tau_{m,k}}^{t-1}(1 + \mathcal{R}^{DOM}_{f,w:w+1})$. If no break has been detected before or at t, such that $k = 0$, then $s^B_{t,DOM} = s^A_{t,DOM} = s^A_{[mT]+1,DOM} \cdot \prod_{w=0}^{t-1}(1 + \mathcal{R}^{DOM}_{f,w:w+1})$.

For example, let USD be the domestic currency, i.e. $DOM =$ USD. From Table 3, the triplet JPY-XBT-EUR experiences a break-point on 29 January 2014 (that is $\tau_{m,1} = [mT] + 52$). On this date, the spot rates were $S_{\tau_{m,1},USD/JPY} = 0.0098$, $S_{\tau_{m,1},USD/XBT} = 834.9783$ and $S_{\tau_{m,1},EUR/USD} = 0.7320$.

Given the daily yen-denominated risk-free rates, the initial units of 905.8119 JPY will be worth 905.9270 JPY at time $\tau_{m,1}$. The trading strategy will exchange the 905.9270 units of JPY for $905.9270 \times S_{\tau_{m,1},USD/JPY} = 8.8659$ units of USD. Given initial spot rates $S_{[mT]+1,XBT/USD} = 0.0027$ and $S_{[mT]+1,EUR/USD} = 0.7447$, the strategy will also convert the units of XBT and EUR into 20.4128

and 9.2517 units of USD.[24] The return for Portfolio B is defined as

$$\mathcal{R}^{B}_{[mT]+1:T} = \frac{\left(\sum_{i=1}^{11} s^{B}_{T,i} S_{T,DOM/i} - 100\right)}{100} \ , \tag{8}$$

where $s^{B}_{T,i}$ are the units of currency i owned on the liquidation date T. In addition to the portfolio returns $\mathcal{R}^{A}_{[mT]+1:T}$ and $\mathcal{R}^{B}_{[mT]+1:T}$, we also calculate a risk-adjusted measure of portfolio performance, i.e., the Sharpe Ratio for portfolio $\mathcal{P} \in \{A, B\}$, as the ratio of mean daily portfolio excess returns to the standard deviation of portfolio daily returns:

$$SharpeRatio^{\mathcal{P}} = \sqrt{T - [mT] - 1} \frac{\frac{1}{T-[mT]-1} \sum_{t=[mT]+1}^{T} \left(\mathcal{R}^{\mathcal{P}}_{t-1:t} - \mathcal{R}^{DOM}_{f,t-1:t}\right)}{\hat{\sigma}(R^{\mathcal{P}}_{t-1:t})} \ , \tag{9}$$

where $\hat{\sigma}(\cdot)$ denotes sample standard deviation of the daily portfolio returns

$$\mathcal{R}^{\mathcal{P}}_{t-1:t} := \frac{\sum_{i=1}^{11} \left(s^{\mathcal{P}}_{t,i} S_{t,DOM/i} - s^{\mathcal{P}}_{t-1,i} S_{t-1,DOM/i}\right)}{\sum_{i=1}^{11} s^{\mathcal{P}}_{t-1,i} S_{t-1,DOM/i}}$$

and $\mathcal{R}^{DOM}_{f,t-1:t}$ the daily "risk-free" rate of the domestic currency. Again, for the risk-free rate, we use the one-month deposit rate of the respective domestic currency, as obtained from Thomson Reuters Datastream.

Table 4 compares the buy-and-hold returns and Sharpe ratios of the two portfolios, for the different domestic currencies USD (Panel I), EUR (Panel II), CHF (Panel III), and JPY (Panel IV). In each of the panels, Portfolio B is shown to significantly outperform Portfolio A, both in terms of returns and in terms of the Sharpe ratio. In fact, Panel I and Panel III show that, when the USD or CHF is the domestic currency, the buy-and-hold Portfolio A earns a strongly negative return $\mathcal{R}^{A}_{[mT]+1:T}$, while Portfolio B achieves a positive return.[25]

[Insert Table 4 about here.]

It could be that the higher profits of Portfolio B simply reflect general movements and appreciations

[24] Note that, once a foreign currency is sold, we do not allow the trader to buy back the currencies afterwards. A trading strategy based on an econometric tool would require to estimate the particular duration of a financial market dislocation, for example by monitoring the triangular arbitrage parity deviations for a break from no cointegration to cointegration. We leave this extension for future research.

[25] We perform an additional robustness check in which the portfolio strategies considered without any interest payments, i.e., holding pure cash portfolios. The result that Portfolio B outperforms Portfolio A both in terms of returns and Sharpe ratios remains qualitatively unchanged.

in currency markets, rather than a result of optimally-timed market trades. Therefore, as a robustness check, we next perform a stability analysis in which we compare the performance of our Portfolio B with a portfolio that trades on randomly chosen break-point dates. In these "placebo" portfolios, the number of break-points is chosen from a Poisson distribution with mean equal to the observed number of break-points, $\lambda = 15$. Break-point dates and are chosen from a uniform distribution on the grid $[mT] + 1, \ldots, T$, and associated currency triplets are chosen from a uniform distribution without replacement. For the four domestic currencies considered, the portfolio are simulated a total of 5,000 times; the distributions of resulting Sharpe ratios are shown in Figure 11. In each figure, the red vertical line corresponds to the Sharpe ratio from the portfolio that trades based on the break-points as detected by the Wagner and Wied (2017) monitoring procedure. It is clear that the Portfolios B that apply the Wagner and Wied (2017) monitoring procedure significantly outperform those portfolios that trade based on randomly assigned break-points, further strengthening the argument that the Wagner and Wied (2017) procedure indeed detects periods of instability in which it is optimal to exit currency markets.

[Insert Figure 11 about here.]

Overall, this exercise is meant as a simple illustration that our methodology could lead to portfolio strategies that outperform, e.g., a simple buy-and-hold strategy. The fact that our portfolio consistently outperforms the buy-and-hold portfolio supports the idea that our monitoring tool indeed detects financial market dislocations, which represent instances of heightened uncertainty in foreign exchange markets. This characteristic of the monitoring tool should allow its use for constructing profitable investment strategies.

6 Summary

This paper uses the monitoring tools of Wagner and Wied (2017) to investigate deviations from the triangular arbitrage parity in foreign exchange markets. In this way, we contribute to the literature on financial market dislocations by applying new econometric tools to detect substantial mispricings in the market. To examine this issue, we collect spot exchange rates for ten fiat currencies, and use transaction data to construct spot rates for the cryptocurrency Bitcoin. Using the stationarity monitoring tool from Wagner and Wied (2017), which monitors the stability of a cointegrating relationship over time, our

results confirm that deviations from the triangular arbitrage parity are rare for fiat currencies. From a total of 72 different triplet permutations of traditional fiat currencies, we find a structural break in the relationship implied by the triangular arbitrage parity in only two cases.

On the other hand, for currency triplets involving the cryptocurrency Bitcoin, we detect break-points in the majority of cases. Most of the detected break-points closely correspond to market events, such as the Swiss National Bank's surprise lifting of its cap on the Swiss franc, and also to the bankruptcy of and the resulting trading halt on Mt. Gox, which at the time of its bankruptcy was the largest Bitcoin-trading platform in the world. This perhaps reflects the finding as in Glaser et al. (2014) that Bitcoin is more a speculative asset than it is a currency.

References

Akram, Q. F., Rime, D., and Sarno, L. (2008). Arbitrage in the foreign exchange market: Turning on the microscope. *Journal of International Economics*, 76(2):237–253.

Andrews, D. W. K. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–854.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866.

Barro, R. J. (2009). Rare disasters, asset prices, and welfare costs. *The American Economic Review*, 99(1):243–264.

Bollerslev, T. and Todorov, V. (2011). Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211.

Campbell, J. Y., Serfaty-De Medeiros, K., and Viceira, L. M. (2010). Global currency hedging. *The Journal of Finance*, 65(1):87–121.

Campbell, J. Y., Viceira, L. M., and White, J. S. (2003). Foreign currency for long-term investors. *The Economic Journal*, 113(486):C1–C25.

Courtois, N. T., Emirdag, P., and Nagy, D. A. (2014). Could Bitcoin transactions be 100x faster? In *Proceedings of the 11th International Conference on Security and Cryptography (SECRYPT).*

Decker, C. and Wattenhofer, R. (2014). Bitcoin transaction malleability and MtGox. In *Proceedings of the European Symposium on Research in Computer Security.*

Dong, H. and Dong, W. (2014). Bitcoin: Exchange rate parity, risk premium, and arbitrage stickiness. *British Journal of Economics, Management & Trade*, 5(1):105–113.

Engel, C. and Rogers, J. (1996). How wide is the border? *American Economic Review*, 86(5):1112–1125.

Evans, M. D. (2018). Forex trading and the WMR fix. *Journal of Banking & Finance*, 87:233–247.

Fink, C. and Johann, T. (2014). Bitcoin markets. Working Paper.

Fleckenstein, M., Longstaff, F. A., and Lustig, H. (2014). The TIPS-Treasury bond puzzle. *The Journal of Finance*, 69(5):2151–2197.

Foley, S., Karlsen, J., and Putniņš, T. J. (2018). Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? Working Paper.

Foucault, T., Kozhan, R., and Tham, W. W. (2016). Toxic arbitrage. *The Review of Financial Studies*, 30(4):1053–1094.

Gagnon, L. and Karolyi, A. (2010). Multi-market trading and arbitrage. *Journal of Financial Economics*, 97(1):53–80.

Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M., and Siering, M. (2014). Bitcoin-asset or currency? Revealing users' hidden intentions. In *Proceedings of the 22nd European Conference on Information Systems*.

Gromb, D. and Vayanos, D. (2010). Limits to arbitrage: The state of the theory. *Annual Review of Financial Economics*, 2:251–275.

Hau, H. and Rey, H. (2006). Exchange rates, equity prices, and capital flows. *The Review of Financial Studies*, 19(1):273–317.

Kim, T. (2017). On the transaction cost of bitcoin. *Finance Research Letters*, 23:300–305.

Klemkosky, R. C. and Lee, J. H. (1991). The intraday ex post and ex ante profitability of index arbitrage. *Journal of Futures Markets*, 11(3):291–311.

Kozhan, R. and Tham, W. W. (2012). Execution risk in high-frequency arbitrage. *Management Science*, 58(11):2131–2149.

Kwiatkowski, D., Phillips, P. C., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1):159 – 178.

Lyons, R. K. and Moore, M. J. (2009). An information approach to international currencies. *Journal of International Economics*, 79(2):211–221.

Makarov, I. and Schoar, A. (2018). Trading and arbitrage in cryptocurrency markets. Working Paper.

Marsh, I. W., Panagiotou, P., and Payne, R. (2017). The WMR fix and its impact on currency markets. Working Paper.

Matvos, G. and Seru, A. (2014). Resource allocation within firms and financial market dislocation: Evidence from diversified conglomerates. *The Review of Financial Studies*, 27(4):1143–1189.

Modest, D. M. and Sundaresan, M. (1983). The relationship between spot and futures prices in stock index futures markets: Some preliminary evidence. *Journal of Futures Markets*, 3(1):15–41.

Moore, T. and Christin, N. (2013). Beware the middleman: Empirical analysis of Bitcoin-exchange risk. In *Proceedings of the International Conference on Financial Cryptography and Data Security.*

Pasquariello, P. (2014). Financial market dislocations. *Review of Financial Studies*, 27(6):1868–1914.

Phillips, P. C. B. and Hansen, B. E. (1990). Statistical inference in instrumental variables regression with I(1) processes. *Review of Economic Studies*, 57(1):99–125.

Pieters, G. and Vivanco, S. (2017). Financial regulations and price inconsistencies across Bitcoin markets. *Information Economics and Policy*, 39:1–14.

Pieters, G. C. (2016). Does bitcoin reveal new information about exchange rates and financial integration? Working Paper.

Saikkonen, P. (1991). Asymptotically efficient estimation of cointegration regressions. *Econometric Theory*, 7(1):1–21.

Shleifer, A. and Vishny, R. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Skiadas, C. (2009). *Asset Pricing Theory.* Princeton Series in Finance. Princeton University Press.

Stock, J. H. and Watson, M. W. (1993). A simple estimator of cointegrating vectors in higher order integrated systems. *Econometrica*, 61(4):783–820.

Veronesi, P. (2004). The peso problem hypothesis and stock market returns. *Journal of Economic Dynamics and Control*, 28(4):707–725.

Vogelsang, T. J. and Wagner, M. (2014). Integrated modified OLS estimation and fixed-b inference for cointegrating regressions. *Journal of Econometrics*, 178(2):741–760.

Wagner, M. and Wied, D. (2017). Consistent monitoring of cointegrating relationships: The US housing market and the subprime crisis. *Journal of Time Series Analysis*, 38(6):960–980.

Wied, D., Arnold, M., Bissantz, N., and Ziggel, D. (2012). A new fluctuation test for constant variances with applications to finance. *Metrika*, 75(8):1111–1127.

Wied, D., Ziggel, D., and Berens, T. (2013). On the application of new tests for structural changes on global minimum-variance portfolios. *Statistical Papers*, 54(4):955–975.

Yu, G. and Zhang, J. (2017). Revisit capital controls policies when bitcoin is in town. Working Paper.

A Appendix

A.1 Times Series Properties of Exchange Rate Data

Table 5 presents ISO 4217 Currency Codes, symbols, and share of trading for the fiat currencies used in this analysis. Liquidity of the currencies ranges from the U.S. dollar (USD), which is involved in 87.6% of currency transactions, to the South African Rand (ZAR), which is involved in only 1% of all currency transactions.[26]

[Insert Table 5 about here.]

Exchange rate quotations A/B are expressed as the units of currency A received for one unit of currency B. For stationarity and cointegration monitoring, one central assumption in our analysis is that log exchange rates are $I(1)$. Therefore, we test each exchange rate pair individually for $I(1)$ behavior using common tests for stationarity and present the results in Table 6, for each currency paired against the USD.[27] Panel A shows that, as expected, the augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests fail to reject the null hypothesis of a unit root when testing all log exchange rates in levels. Likewise, the Kwiatkowski et al. (1992) (KPSS) test rejects the null of stationarity for all log exchange rates. On the other hand, the test results from Panel B shows that first-differences of all log exchange rates exhibit stationarity. Therefore, evidence from these tests affirms that log exchange rates are indeed $I(1)$.[28]

[Insert Table 6 about here.]

A.2 Precision of Bitcoin Data

Table 7 shows summary statistics for the precision of our Bitcoin noon exchange rates with respect to the target window of 11:59am and 12:01pm ET. As expected, USD/XBT and EUR/XBT are the most precise: 88% and 65% of our noon exchange rates capture only observations within the target window, and typically more than three observations are found within the window (an average of 65 observations

[26] Note that, as multiple currencies are involved in a currency trade, the sum of the average daily turnover rates for individual currencies is greater than 100%.

[27] Results for currencies paired against the EUR are similar.

[28] In addition, to exclude the possibility of cointegration relationships between the components of $\mathbf{x}_t$, we run Johansen cointegration tests on pairs of exchange rates $\mathbf{x}_t = \left(\ln S_{t,A/V}, \ln S_{t,V/B}\right)^\top$. The tests assess the null hypothesis of a cointegration rank $\leq r$, where $r = 0$ or $r = 1$. We fail to reject the null for 89% of the currency pairs in the case of $r = 0$, and fail to reject the null for 99% of the currency pairs in the case of $r = 1$.

for USD/XBT, and 13 for EUR/XBT). The Bitcoin rate against the Mexican peso, MXN/XBT, is the least precise: results show that a MXN/XBT observation never occurs during the target window, and the median time distance of observations away from the target window is about ten hours. All-in-all, the results show that our methodology captures rates that on average occur at least within a 12-hour window around the target window. Figure 12 plots the histogram of observations used to calculate Bitcoin noon exchange rates that fall outside of the target window. Observations are binned according to hours away from the target window. As expected, USD/XBT and EUR/XBT are clustered around 0, while other currencies have wider tails. Interestingly, most currencies see a spike around 12 hours after the target window; this might be because of Bitcoin exchange operating hours, or due to time zone effects of the countries in which the paired fiat currency is traded.

[Insert Table 7 and Figure 12 about here.]

We believe that this approach is superior to using daily averages of Bitcoin prices. Using only the Bitcoin prices that are within or closest to the same two-minute window as our fiat currency spot rates allows us to more accurately capture prices that would be present in the market at approximately the same time, and thus be available simultaneously to traders looking to engage in triangular arbitrage. However, there are a few drawbacks. First, given the 24-hour nature of Bitcoin exchange trading and the dominance of non-U.S. exchanges (Fink and Johann (2014) find that, of the top five most liquid Bitcoin exchanges, none are based in the United States), a 11:59am-12:01pm ET time window may not necessarily be a representative time period for Bitcoin markets. Secondly, Bitcoin transactions have a much higher latency than transactions in fiat currencies, due to the need to verify and confirm Bitcoin transactions within the blockchain network (see, e.g., Courtois et al., 2014). Therefore, traders may only be able to exploit arbitrage opportunities with some amount of lag in Bitcoin markets that is longer than the two-minute window.

A.3 Bid-Ask Spreads

We defined the percentage bid-ask spread as the relative bid-ask spread in percentage terms $\delta_{t,A/B} = \frac{\kappa_{t,A/B}}{S^a_{t,A/B}} \times 100$.[29] Let $\breve{\delta}_{t,A/B} = \frac{\kappa_{t,A/B}}{S^a_{t,A/B}}$. Note that, since $S^b_{t,A/B} \leq S_{t,A/B} \leq S^a_{t,A/B}$, we have $S^a_{t,A/B} \leq$

[29] We calculate percentage bid-ask spreads as $\delta_{t,A/B} = \frac{\kappa_{t,A/B}}{S^a_{t,A/B}} \times 100$, as this is how spreads are calculated by the data provider Bitcoinity, which does not provide data on the individual bid and ask quotes. Dividing by the larger price therefore may serve to underestimate the percentage spread. However, we do not believe that we underestimate Bitcoin spreads for

$S_{t,A/B} + \kappa_{t,A/B} = S_{t,A/B}\left(1 + \breve{\delta}_{t,A/B}\right)$ and $S^b_{t,A/B} \geq S_{t,A/B} - \kappa_{t,A/B} = S_{t,A/B}\left(1 - \breve{\delta}_{t,A/B}\right)$. By taking logarithms we observe that:

$$\begin{aligned}
\ln S^a_{t,A/B} \quad &\leq \ln S_{t,A/B} + \ln\left(1 + \breve{\delta}_{t,A/B}\right) \leq \ln S_{t,A/B} + \ln\left(1 + \max_t \breve{\delta}_{t,A/B}\right) \\
&\leq \ln S_{t,A/B} + \left|\ln\left(1 - \max_t \breve{\delta}_{t,A/B}\right)\right| \ , \text{ and} \\
\ln S^b_{t,A/B} \quad &\geq \ln S_{t,A/B} + \ln\left(1 - \breve{\delta}_{t,A/B}\right) \geq \ln S_{t,A/B} + \ln\left(1 - \max_t \breve{\delta}_{t,A/B}\right) \ .
\end{aligned}$$

Hence, an upper bound on trading costs implied by the spread follows from $\bar{s}_{A/B} := \max_t |\ln(1 - \breve{\delta}_{t,A/B})| = \max_t |\ln(1 - \delta_{t,A/B}/100)|$. In addition, since the Bitcoin bid-ask spreads are obtained from various exchanges and include some extreme realizations which directly enter into our $\bar{s}$, to exclude possible effects arising from such spikes in the bid-ask spread, we additionally replace $\max_t$ by the 90% percentile and obtain $\bar{s}_{A/B,90\%}$. To obtain a lower bound and estimates of the mean location of the spread, we make the simplifying assumption that $S_{t,A/B}$ is the midpoint between $S^a_{t,A/B}$ and $S^b_{t,A/B}$. This implies that traders on average only pay the half-spread, yielding an implied transaction cost of $\tilde{\delta}_{t,A/B} := \frac{\kappa_{t,A/B}}{2S_{t,A/B}} \times 100$. By considering the minimum, the mean and the median of $|\ln(1 - \tilde{\delta}_{t,A/B}/100)|$ we obtain a lower bound $\underline{s}_{A/B} = \min_t |\ln(1 - \tilde{\delta}_{t,A/B}/100)|$, a measure of the mean relative spread $s_{A/B,mean}$, and a measure of the median relative spread $s_{A/B,median}$. This procedure is applied to all exchange rates. Taking the sums, e.g. $\bar{s} := \bar{s}_{A/B} + \bar{s}_{A/V} + \bar{s}_{V/B}$, provides us with measures of the largest spread $\bar{s}$, the 90% percentile $s_{90\%}$, the mean s_{mean}, the median s_{median} and the smallest spread $\underline{s}$.

several reasons. First, Bitcoin spreads calculated by Bitcoinity are calculated as the difference between the minimum bid and maximum ask, therefore representing a daily upper bound. Secondly, we take the maximum across exchanges. Our resulting calculations of Bitcoin spreads larger by a magnitude of 10 than those found in, e.g., Kim (2017). Meanwhile, underestimating fiat currency bid-ask spreads would only bias our results towards finding more deviations outside of spread-implied no-arbitrage bounds, which we do not find. Therefore we do not believe that our procedure adversely underestimates the relative bid-ask spread in calculating the lower and upper bounds.

A.4 Cointegration Monitoring

A further way to monitor market dislocations is *cointegration monitoring*, which is also proposed in Wagner and Wied (2017). Given the time series properties of exchange rates, the process $\{\mathbf{x}_t\}_{t\in\mathbb{Z}}$, $\mathbf{x}_t := \left(\ln S_{t,A/V}, \ln S_{t,V/B}\right)^\top$, follows a vector random walk process $\mathbf{x}_t = \mathbf{x}_{t-1} + \mathbf{v}_t$, where $\{\mathbf{v}_t\}_{t\in\mathbb{Z}}$ is integrated of order zero ($I(0)$). The stacked process $\{\boldsymbol{\eta}_t\}_{t\in\mathbb{Z}}$, $\boldsymbol{\eta}_t = (u_t, \mathbf{v}_t^\top)^\top$ is $I(0)$ under the null hypothesis. To obtain the asymptotic limit distribution, the authors assume that a functional central limit theorem holds. That is, $\frac{1}{\sqrt{T}}\sum_{t=1}^{[rT]} \boldsymbol{\eta}_t$ weakly converges to the Brownian motion $\mathbf{B}(r) = \boldsymbol{\Omega}^{1/2}\mathbf{W}(r)$, where $\boldsymbol{\Omega}$ denotes the long-run covariance matrix of $\boldsymbol{\eta}_t$ and $\mathbf{W}(r)$ is a standard Brownian motion $\in \mathbb{R}^3$. Note that:

$$\boldsymbol{\Omega} = \begin{bmatrix} \Omega_{uu} & \boldsymbol{\Omega}_{uv} \\ \boldsymbol{\Omega}_{vu} & \boldsymbol{\Omega}_{vv} \end{bmatrix} = \begin{bmatrix} \omega^2 & \boldsymbol{\Omega}_{uv} \\ \boldsymbol{\Omega}_{vu} & \boldsymbol{\Omega}_{vv} \end{bmatrix} := \sum_{j=-\infty}^{\infty} \mathbb{E}(\boldsymbol{\eta}_{t-j}\boldsymbol{\eta}_t^T) \ ,$$

where $\boldsymbol{\Omega}$ is assumed to be regular and finite. By this assumption, cointegration relationships between the components of $\mathbf{x}_t$ are excluded. Given these assumptions, equation (2) describes a cointegrating relationship.

The observations $(y_t,\ \mathbf{x}_t :\ 1 \le t \le [mT])$ are used to estimate α and $\boldsymbol{\beta}$ using fully modified least squares (FM-OLS; see, e.g. Phillips and Hansen, 1990), dynamic least squares (DOLS; see, e.g., Saikkonen, 1991; Stock and Watson, 1993) or integrated least squares (IM-OLS; see, e.g., Vogelsang and Wagner, 2014). Denoting the FM-residuals by $\widehat{u}_{t,m}^+$, where $\widehat{u}_t^+ = y_t - \widehat{\alpha}_m^+\mathbf{D}_t - \widehat{\boldsymbol{\beta}}_m^{+\top}\mathbf{x}_t$, FM-residuals $\widehat{u}_t^+$ are then used to obtain the test statistic $\widehat{H}^{m+}(s)$.[30] α and $\boldsymbol{\beta}$ are estimated with FM-OLS, and the asymptotic limit distribution of $\widehat{H}^{m+}(s)$ depends on the estimation methodology, as well as on the choice of m, $\dim(\mathbf{x}_t)$, and whether equation (2) includes a constant and/or linear trend. For DOLS and IM-OLS we can proceed in a similar way. For more details we refer the reader to Wagner and Wied (2017).

In addition, estimating α and $\boldsymbol{\beta}$ provides us with the opportunity to run a Wald-type test to infer whether $(\alpha, \boldsymbol{\beta}^\top)^\top = (\alpha^*, \boldsymbol{\beta}^{*\top})^\top = (0,1,1)^\top$, and separately whether $\alpha = \alpha^* = 0$ and $\boldsymbol{\beta} = \boldsymbol{\beta}^* = (1,1)^\top$. Let $\widehat{\alpha}_m^+$ and $\widehat{\boldsymbol{\beta}}_m^+$ denote the FM-OLS estimates of the parameters α and $\boldsymbol{\beta}$ based on the observations $t = 1, \ldots, [mT]$. By means of these estimates we test whether deviations of the parameters from the values implied by the triangular arbitrage parity can be observed in our dataset.

Hence, by assuming that the log-exchange rate time series are integrated of order one ($I(1)$), we

[30] The asymptotic limit distributions for FM-OLS, DOLS, and IM-OLS are provided in Wagner and Wied (2017). $\widehat{H}^{m+}(s)$ is obtained in a similar way as $\widehat{H}^m(s)$ defined in (4).

can run *cointegration monitoring* and test whether the estimated parameters are significantly different from the theory-implied values. We analyze three Wald-type tests on the estimated parameters: first, a joint test of all parameters, $\alpha = 0$ and $\beta_1 = \beta_2 = 1$ ("all-parameters hypothesis"); a test of $\alpha = 0$ ("alpha hypothesis"), and a joint test of the betas, $\beta_1 = \beta_2 = 1$ ("beta hypothesis").

For the USD triplets, the Wald-type test rejects the null hypothesis that all parameters are equal to their theory-implied values for one-third of the fiat-only currency triplets, but for all of the currency triplets that include Bitcoin. We reject the "alpha null hypothesis" that $\alpha = 0$ in five (14%) of the fiat-only triplets, and in five (56%) of the Bitcoin triplets. Likewise, the "beta test" rejects the "beta null hypothesis" that $\beta_1 = \beta_2 = 1$ in three (8%) of the fiat-only triplets, and in 3 (33%) of the Bitcoin triplets. Likewise, for EUR triplets, we reject the "all-parameters null hypothesis" for six (17%) of the fiat-only and all of the Bitcoin triplets, the "alpha hypothesis" for five (14%) of the fiat-only and five (56%) of the Bitcoin triplets, and lastly the "beta hypothesis" is rejected for just two (5.5%) and three (33%) of the fiat-only and Bitcoin triplets, respectively.

B Tables and Figures

Figure 1: Time Variation in Exchange Rates

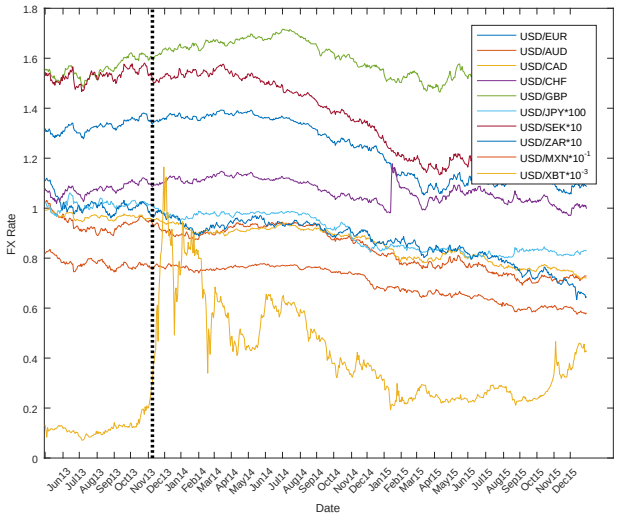

This figure plots exchange rates for ten currencies against the U.S. dollar (USD), during the time period from 1 May 2013 – 31 December 2015. The black dotted line corresponds to the cut-off between the calibration time period and the monitoring time period.

Figure 2: Percentage Bid-Ask Spreads, Fiat Currencies vs. Bitcoin

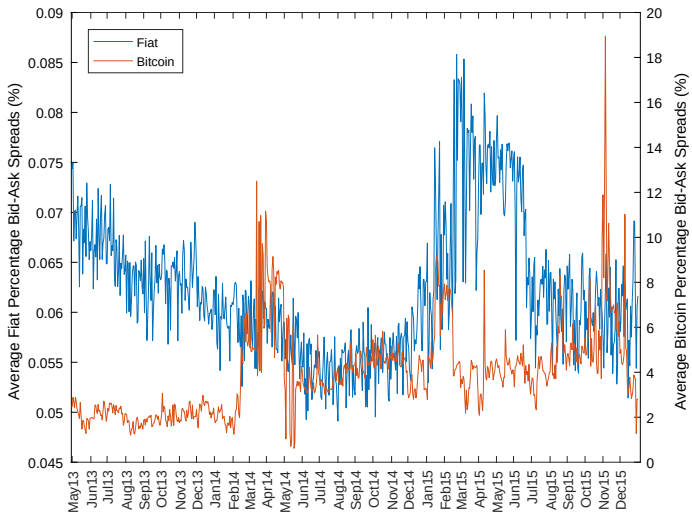

This figure plots the average percentage bid-ask spread for fiat currencies (left axis), as well as average percentage bid-ask spreads for Bitcoin (right axis), during the time period from 1 May 2013 – 31 December 2015.

Figure 3: Deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ for USD-Fiat Currency Triplets

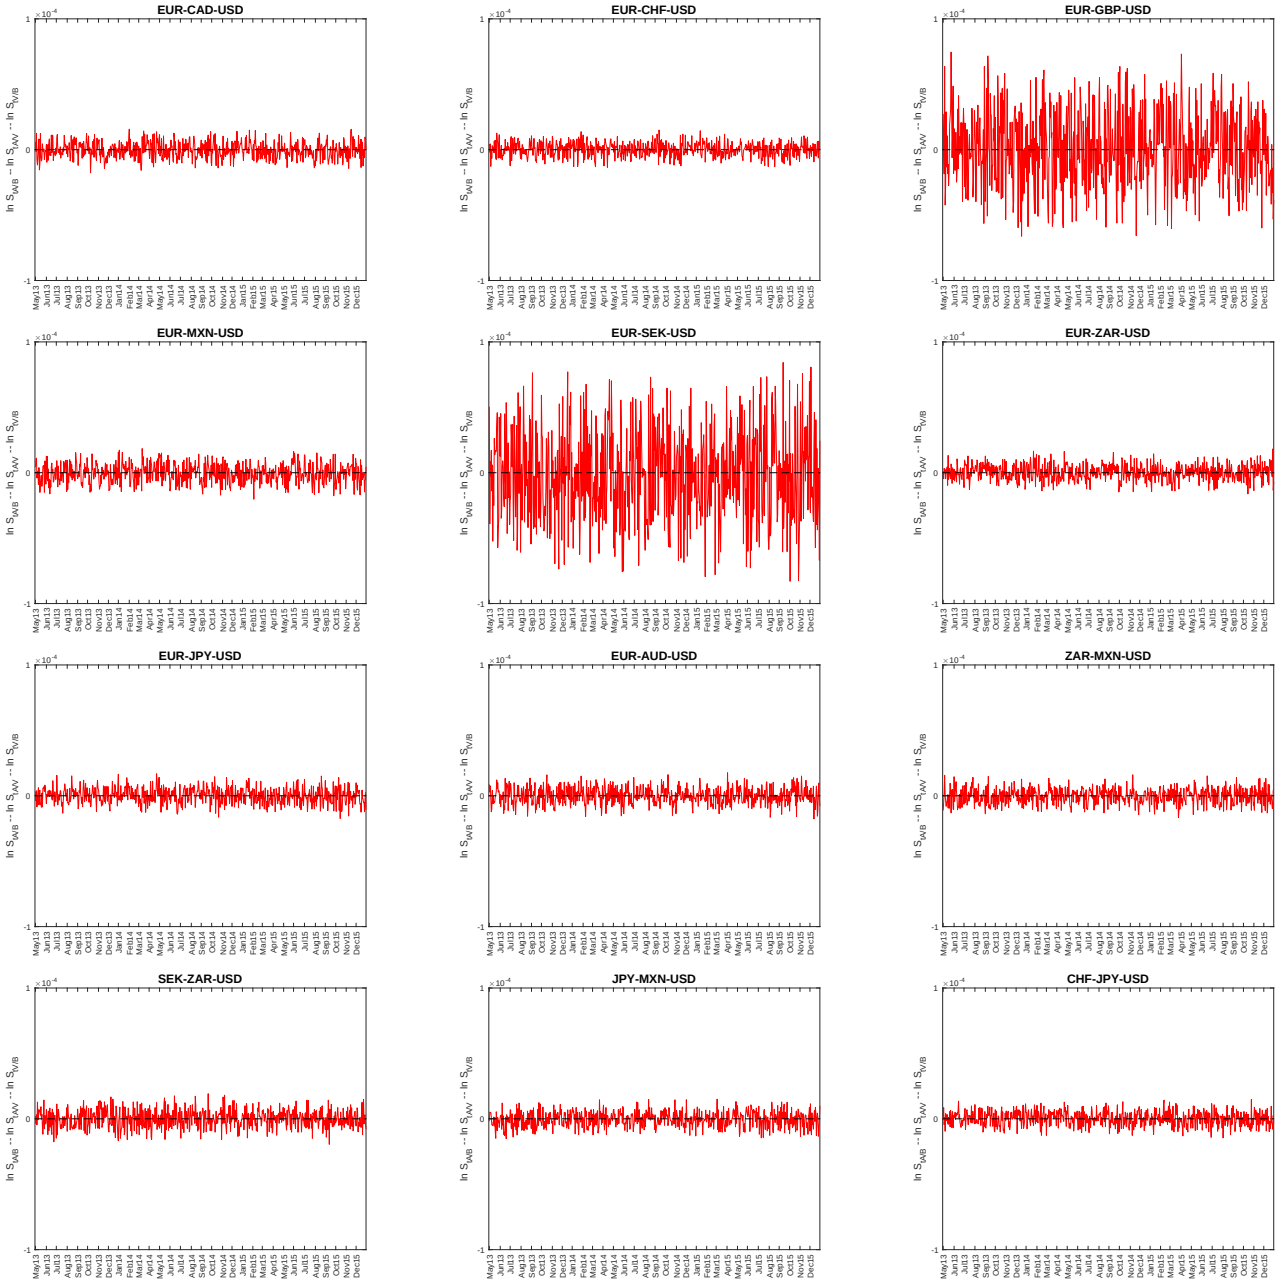

These subfigures plot the deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ for a variety of currency triplets that only include fiat currencies, in which the U.S. dollar (USD) is used as the vehicle currency. The black dotted line corresponds to the cut-off between the calibration and monitoring time periods. The range of the vertical axis is the interval $[-0.0001, 0.0001]$.

Figure 4: Deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ for USD-Bitcoin Triplets

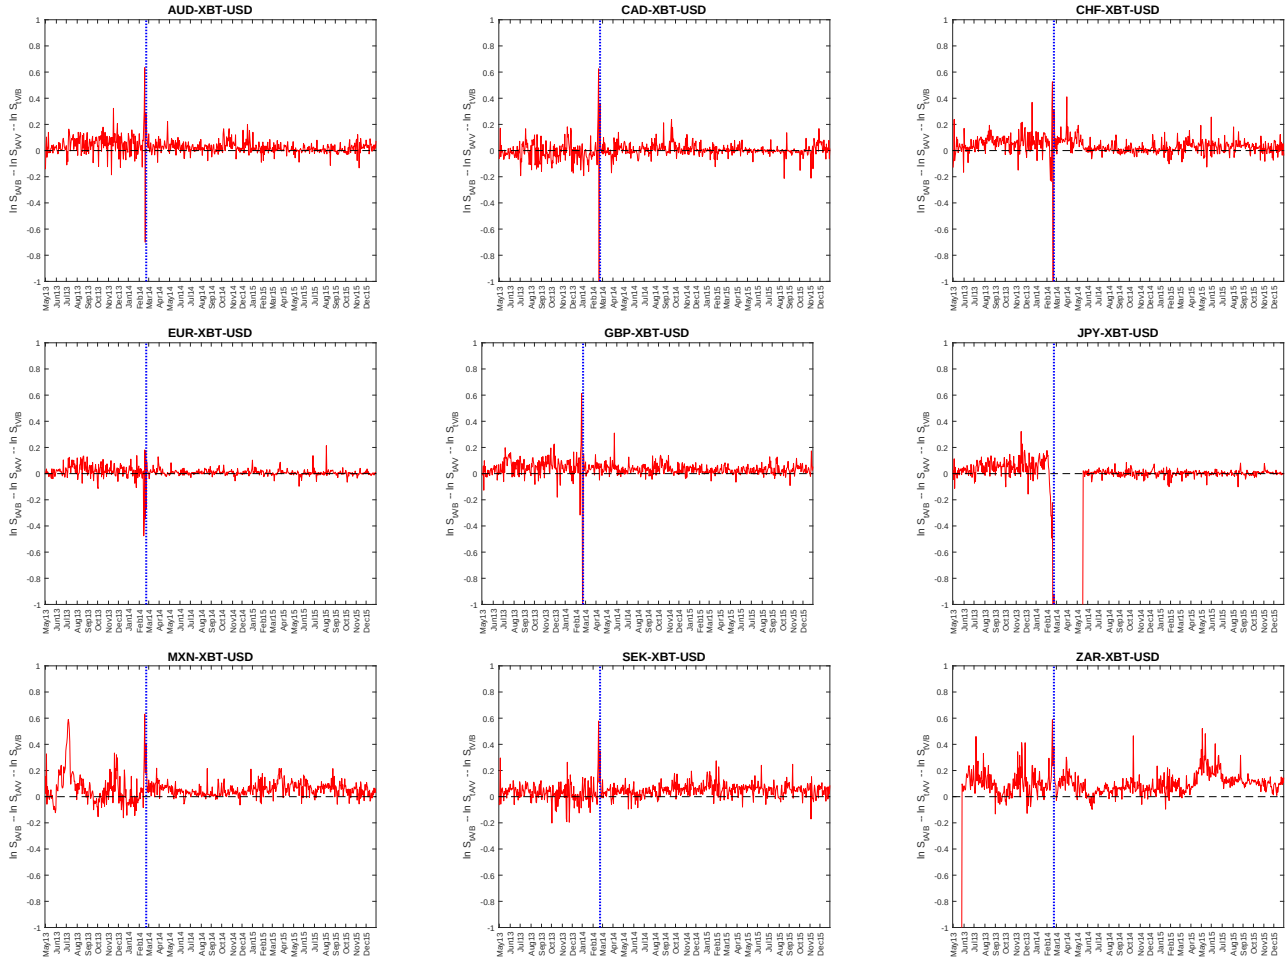

These subfigures plot the deviations $\ln S_{t,A/B} - \ln S_{t,A/V} - \ln S_{t,V/A}$ for a variety of currency triplets that include the crytocurrency Bitcoin (XBT), in which the U.S. dollar (USD) is used as the vehicle currency. The black dotted line corresponds to the cut-off between the calibration and monitoring time periods, while the blue dotted line corresponds to the bankruptcy of Mt. Gox on 24 February 2014. The range of the vertical axis is the interval $[-1, 1]$.

Figure 5: Graphical Illustration of Monitoring Procedure

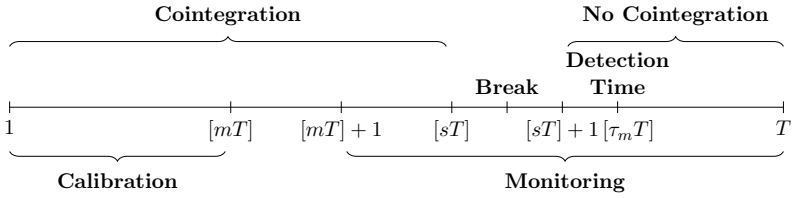

This figure shows a graphical representation of the monitoring procedure from Wagner and Wied (2017). $[mT]$ denotes the next smallest integer of mT, where the time period $t = 1, \ldots [mT]$ denotes the calibration time period. The time period $t = [mT] + 1, ...T$ denotes the monitoring time period, during which a test statistic $\widehat{H}^m(s)$ is constructed for each timepoint $s \in (m, 1]$. The figure illustrates a detected break-point from cointegration to no-cointegration at time $t_{\tau_m} = [\tau_m T]$, i.e., the first s at which $\widehat{H}^m(s)$ exceeds a critical value.

Figure 6: Rolling Window Regressions

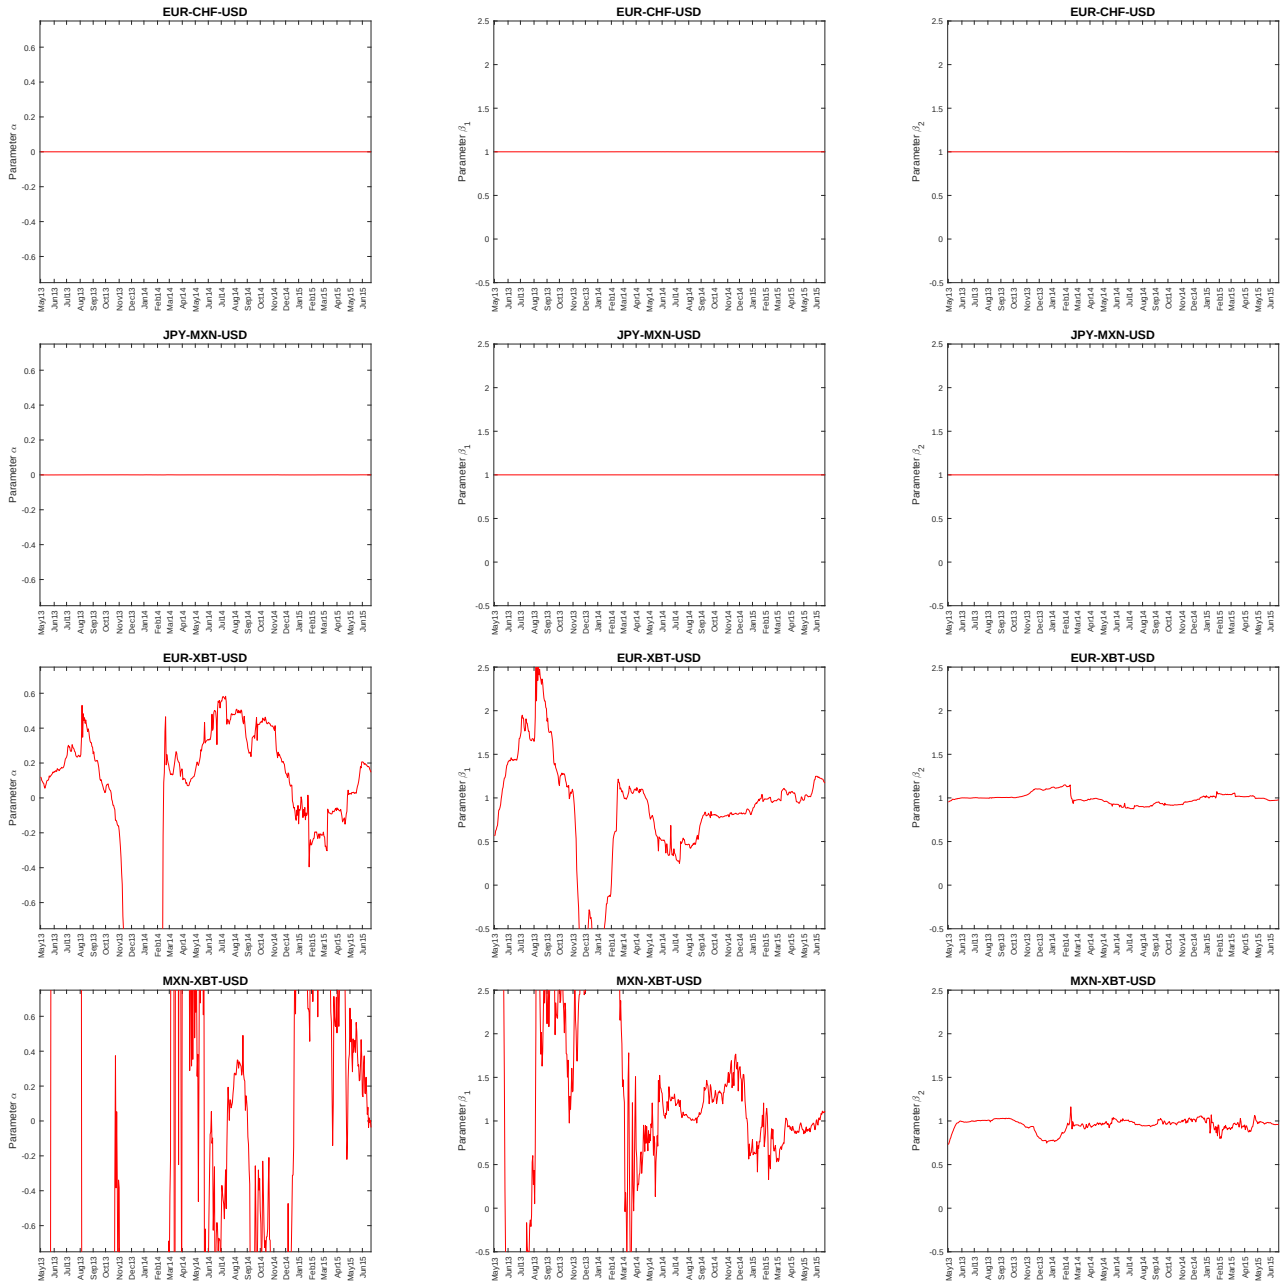

These subfigures show fully modified OLS parameters estimates from rolling regressions in the form of $\ln S_{t,A/B} = \alpha + \beta_1 \ln S_{t,A/V} + \beta_2 \ln S_{t,V/B} + u_t$. The full time period corresponds to 1 May 2013 to 31 December 2015, and uses a time window of $m = 0.2$ to estimate such regressions for $[mT]$ periods. The first column presents estimates of the constant α, the second column for β_1, and the third column for β_2. The triangular arbitrage parity implies the values $\alpha = 0$, and $\beta_1 = \beta_2 = 1$.

Figure 7: Break-Point Dates: USD Triplets

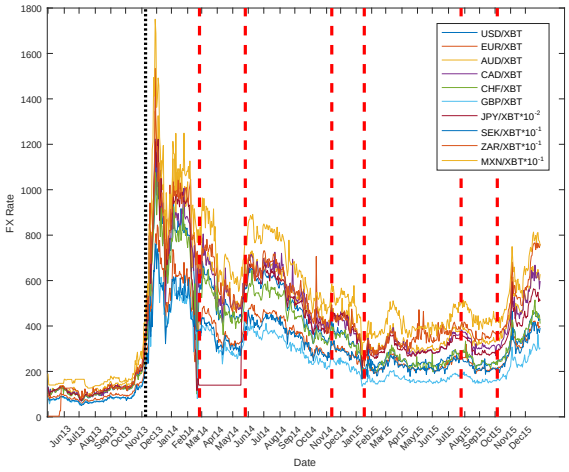

Figure 8: Break-points Dates: EUR Triplets

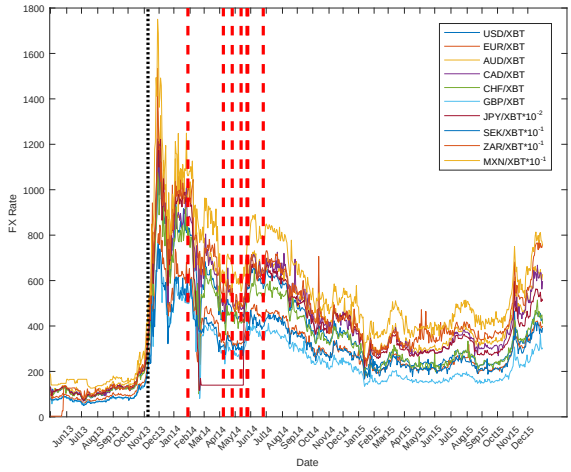

This figure plots exchange rates for nine currencies against Bitcoin (XBT), during the time period from 1 May 2013 – 31 December 2015. The black dotted line corresponds to the cut-off between the calibration time period and the monitoring time period. The red line corresponds to a detected break-point (if any) in the deviations from triangular arbitrage parity for a currency triplet. Figure 7 plots break-points for currency triplets in which the USD is used as a vehicle currency; Figure 8 plots break-points for currency triplets in which the EUR is used as a vehicle currency.

Figure 9: Residuals $\widehat{u}_t$: JPY-XBT-USD

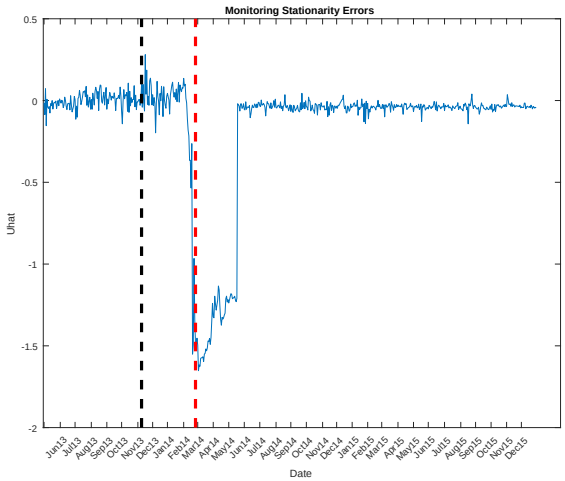

Figure 10: Residuals $\widehat{u}_t$: JPY-SEK-USD

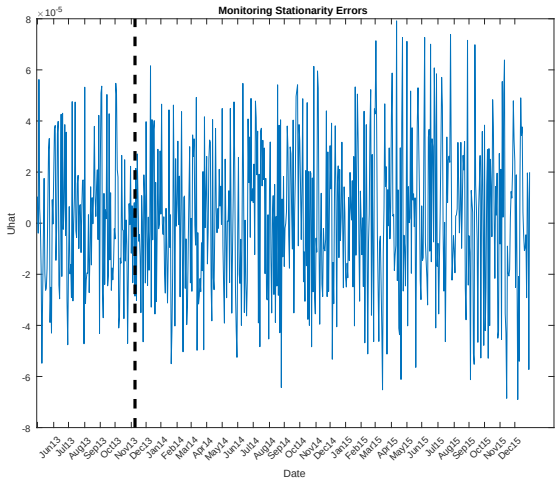

These subfigures plot the OLS residuals $\widehat{u}_t := y_t - \widehat{\alpha}_m - \boldsymbol{\beta}^{*\top}\mathbf{x}_t$ as described in Section 4. The black dotted line corresponds to the cut-off between the calibration and monitoring time periods, and the red dotted line corresponds to the detected break-point (if any) in the deviations from triangular arbitrage parity for the currency triplet shown

Figure 11: Placebo Tests: Portfolio Performance

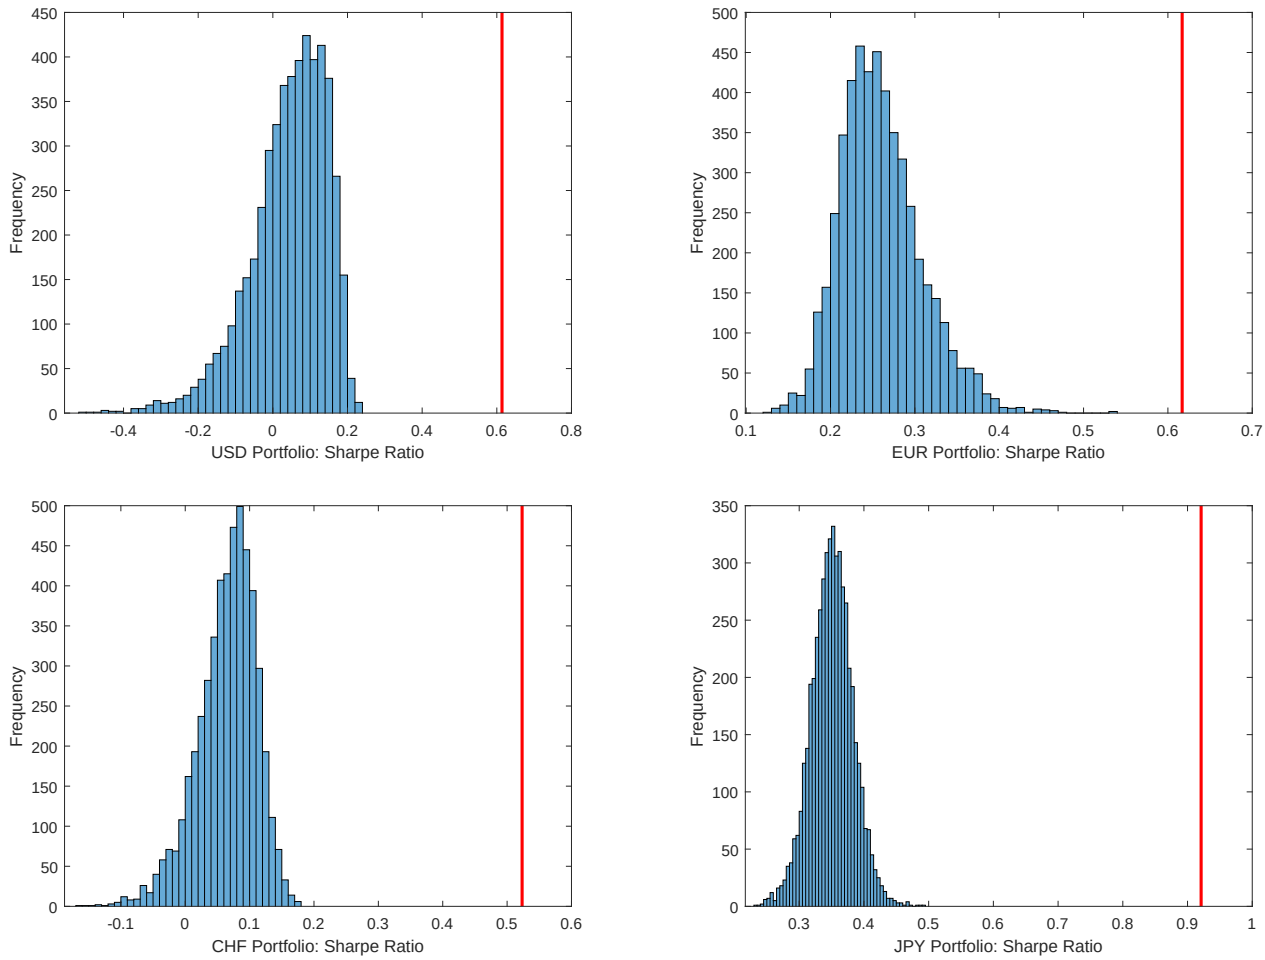

These figures show the distribution of portfolio Sharpe Ratios calculated using randomly assigned break-point dates. The number of break-points is randomly drawn from a Poisson distribution with $\lambda = 15$. Break-point dates are chosen from a uniform distribution. The currency triplet assigned to each break-point date are drawn from a uniform distribution without replacement. Red vertical lines correspond to portfolio Sharpe Ratios calculated using break-points as detected by the Wagner and Wied (2017) monitoring procedure.

Figure 12: Histograms of Bitcoin Observations

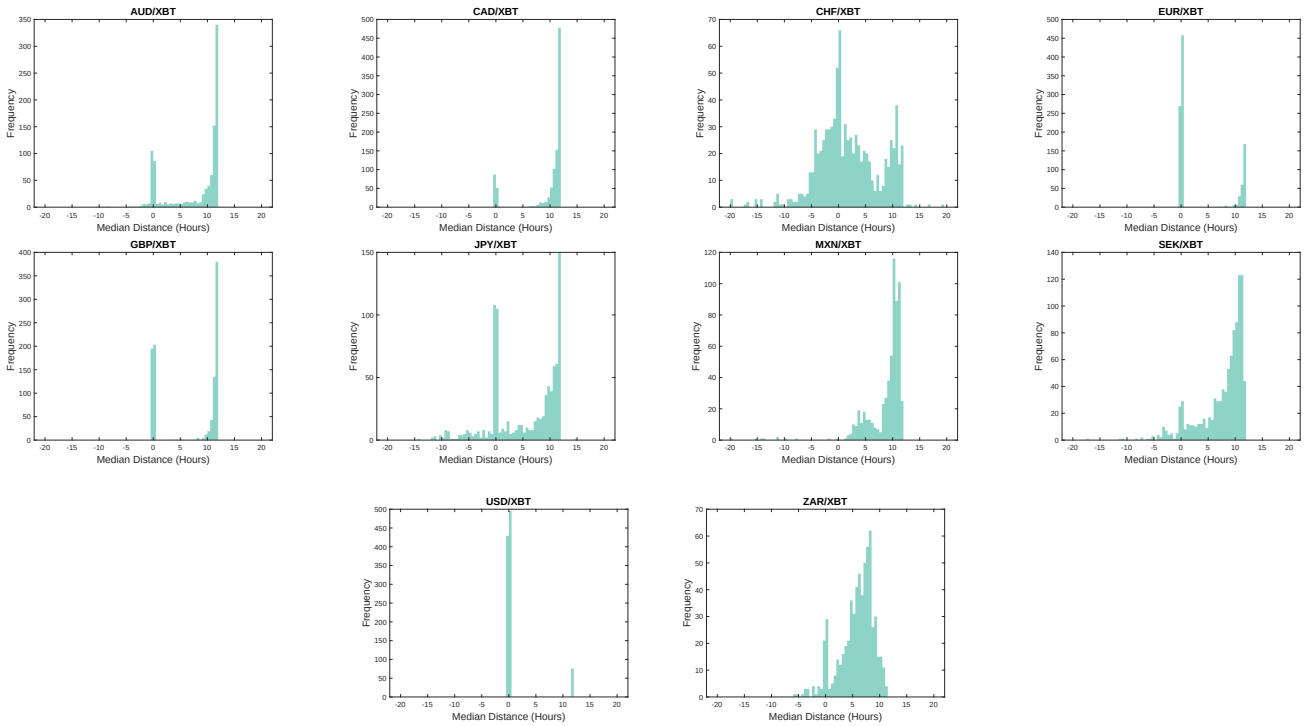

These figures plot the histogram of observations used to calculate Bitcoin noon exchange rates that fall outside of (before or after) the target window of 11:59am to 12:01pm. Histograms are presented separately for each currency.

Table 1: Summary Statistics of Percentage Bid-Ask Spreads

(A) Fiat Currency Percentage Bid-Ask Spreads

	JPY/USD	EUR/USD	AUD/USD	GBP/USD	CAD/USD	SEK/USD	CHF/USD	MXN/USD	ZAR/USD
Mean	0.030	0.023	0.043	0.029	0.030	0.058	0.062	0.019	0.090
Min	0.024	0.014	0.031	0.022	0.021	0.047	0.031	0.007	0.064
Q(50)	0.029	0.023	0.042	0.029	0.029	0.057	0.056	0.019	0.092
Q(90)	0.039	0.028	0.055	0.033	0.038	0.065	0.094	0.028	0.101
Max	0.042	0.032	0.058	0.035	0.060	0.130	0.136	0.041	0.113

(B) Bitcoin Percentage Bid-Ask Spreads

	JPY/XBT	EUR/XBT	AUD/XBT	GBP/XBT	CAD/XBT	SEK/XBT	CHF/XBT	USD/XBT
Mean	9.898	1.417	3.474	3.544	4.375	3.718	9.898	0.414
Min	0.57	0.33	0.467	0.588	0.361	0.895	0.57	0.095
Q(50)	12.836	1.163	3.649	2.053	3.686	3.388	12.836	0.329
Q(90)	15.718	2.31	5.216	7.152	6.381	4.406	15.718	0.7
Max	22.955	10.11	16.832	34.974	46.371	16.833	22.955	7.079

This table shows summary statistics for percentage bid-ask spread for fiat currency exchange rates against the U.S. dollar (Panel A), and for percentage bid-ask spread for exchange rates against Bitcoin (Panel B), during the time period from 1 May 2013 – 31 December 2015. Included are the mean, minimum, median (Q(50)), 90% percentile (Q(90)), and maximum percentage bid-ask spreads.

Table 2: Deviations and Implied Transactions Costs

Triplet	Fiat	EUR XBT USD	AUD XBT USD	CAD XBT USD	CHF XBT USD	GBP XBT USD	JPY XBT USD	SEK XBT USD	ZAR XBT USD	MXN XBT USD
					Vehicle Currency USD					
$\bar{s}$	0.00%	0.90%	0.90%	0.15%	0.45%	0.30%	9.75%	0.75%	NA	NA
$s_{90\%}$	0.00%	22.04%	27.89%	16.34%	2.40%	15.59%	10.34%	50.52%	NA	NA
$\underline{s}$	0.00%	88.76%	92.65%	92.50%	92.20%	94.30%	85.46%	95.80%	NA	NA
s_{mean}	0.00%	59.82%	63.27%	49.48%	58.62%	69.42%	26.99%	79.31%	NA	NA
s_{median}	0.00%	83.21%	89.96%	88.61%	90.25%	92.35%	82.16%	94.75%	NA	NA
					Vehicle Currency EUR					
Triplet	Fiat	USD XBT EUR	AUD XBT EUR	CAD XBT EUR	CHF XBT EUR	GBP XBT EUR	JPY XBT EUR	SEK XBT EUR	ZAR XBT EUR	MXN XBT EUR
$\bar{s}$	0.00%	0.90%	0.90%	0.15%	0.60%	0.60%	9.30%	0.75%	NA	NA
$s_{90\%}$	0.00%	22.04%	14.39%	12.89%	1.80%	6.75%	10.04%	27.89%	NA	NA
$\underline{s}$	0.00%	88.76%	89.96%	90.55%	89.21%	89.96%	80.81%	93.55%	NA	NA
s_{mean}	0.00%	59.82%	50.07%	47.23%	47.38%	56.22%	21.74%	71.21%	NA	NA
s_{median}	0.00%	70.16%	77.36%	79.91%	81.56%	80.81%	68.07%	88.01%	NA	NA

This table presents the relative frequencies at which the deviations $\ln S_{tA/B} - \ln S_{tA/V} - \ln S_{tV/A}$ exceed transaction costs as estimated by the maximum bounds $\bar{s}$ and $\bar{s}_{90\%}$, minimum bound $\underline{s}$, and mean and median estimates s_{mean} and s_{median}, during the period 1 May 2013–31 December 2015. Results are presented separate for currency triplets that only include fiat currencies, and for the nine currency triplets that include Bitcoin (XBT). NA denotes not available.

Table 3: Detected Break-Points

Vehicle Currency USD		Vehicle Currency EUR	
Triplet	$t_{\tau_m} = [\tau_m T]$	Triplet	$t_{\tau_m} = [\tau_m T]$
EUR-XBT-USD	2014/05/27	USD-XBT-EUR	2014/05/27
AUD-XBT-USD	2015/10/06	AUD-XBT-EUR	2014/04/08
CAD-XBT-USD	2014/11/13	CAD-XBT-EUR	2014/04/25
CHF-XBT-USD	2015/01/19	CHF-XBT-EUR	2014/05/13
		GBP-XBT-EUR	2014/05/26
JPY-XBT-USD	2014/02/26	JPY-XBT-EUR	2014/01/29
SEK-XBT-USD	2015/07/27	SEK-XBT-EUR	2014/06/25
CHF-MXN-USD	2014/10/01		
JPY-MXN-USD	2014/10/30		

This table shows all detected break-points found in our sample of currency triplets. Bartlett kernel used to estimate ω^2, $m = 0.2$ and $p = 1$.

Table 4: Comparison of Portfolio Performances

	(I) Domestic Currency: USD		(II) Domestic Currency: EUR	
	Portfolio A	Portfolio B	Portfolio A	Portfolio B
$\mathcal{R}^{\mathcal{P}}_{[mT]+1:T}$	-13.32%	11.61%	7.41%	9.28%
Sharpe Ratio	-0.6421	0.8406	0.4247	0.7019
	(III) Domestic Currency: CHF		(IV) Domestic Currency: JPY	
	Portfolio A	Portfolio B	Portfolio A	Portfolio B
$\mathcal{R}^{\mathcal{P}}_{[mT]+1:T}$	-6.45%	3.72%	4.56%	16.13%
Sharpe Ratio	-0.0401	0.3254	0.2961	0.9911

This table compares the performances of two difference portfolios: Portfolio A, which is based on a buy-and-hold strategy of eleven difference currencies (USD, EUR, JPY, GDP, AUD, CAD, CHF, SEK, MXN, ZAR, and XBT) between 11 November 2013 and 31 December 2015; and Portfolio B, which holds the same currencies but trades foreign currencies for the domestic one when the Wagner and Wied (2017) monitoring procedure detects a break-point in the triangular arbitrage parity deviations of a triplet of currencies. Presented are the buy-and-hold returns $\mathcal{R}^{\mathcal{P}}_{[mT]+1:T}$ and Sharpe ratios for the two portfolios, for strategies when either the USD (Panel I), EUR (Panel II), CHF (Panel III) or JPY (Panel IV) is considered as the domestic currency.

Table 5: Definitions and Trading Shares of Currencies

Currency	Symbol	% Daily Share
United States Dollar	USD ($)	87.6%
Euro	EUR (€)	31.3%
Japanese Yen	JPY (¥)	21.6%
Pound Sterling	GBP (£)	12.8%
Australian Dollar	AUD ($)	6.9%
Canadian Dollar	CAD ($)	5.1%
Swiss Franc	CHF (Fr)	4.8%
Swedish Krona	SEK (kr)	2.2%
Mexican Peso	MXN ($)	1.9%
South African Rand	ZAR (R)	1.0%

This table lists the currencies used in this analysis, along with their ISO 4217 Currency Codes, common symbols (in parenthesis), and percentage share of average daily turnover. The data represents currency turnovers from April 2016 and is sourced from the Bank for International Settlement's Triennial Central Bank Survey (see `http://www.bis.org/publ/rpfx16fx.pdf`). Note that, as multiple currencies are involved in a transaction, the sum of the average daily turnover rates sum to greater than 100%.

Table 6: I(1) Behavior of Log Exchange Rates

(A) Levels

	(1) ADF PVal	(2) ADF Stat	(3) PP PVal	(4) PP Stat	(5) KPSS PVal	(6) KPSS Stat
USD/EUR	0.931	-0.236	0.931	-0.239	< 0.010	8.007
USD/AUD	0.699	-1.091	0.667	-1.163	< 0.010	8.336
USD/CAD	0.958	0.013	0.971	0.171	< 0.010	8.868
USD/CHF	0.109	-2.532	0.087	-2.632	< 0.010	4.288
USD/GBP	0.685	-1.123	0.721	-1.041	< 0.010	3.188
USD/JPY	0.798	-0.867	0.725	-1.032	< 0.010	8.839
USD/SEK	0.889	-0.492	0.885	-0.517	< 0.010	8.835
USD/ZAR	0.976	0.252	0.973	0.213	< 0.010	8.433
USD/MXN	0.936	-0.194	0.947	-0.106	< 0.010	8.537
USD/XBT	0.339	-1.906	0.457	-1.639	< 0.010	1.489

(B) First Differences

	(1) ADF PVal	(2) ADF Stat	(3) PP PVal	(4) PP Stat	(5) KPSS PVal	(6) KPSS Stat
USD/EUR	< 0.001	-9.679	< 0.001	-26.357	> 0.100	0.205
USD/AUD	< 0.001	-9.712	< 0.001	-28.282	> 0.100	0.089
USD/CAD	< 0.001	-8.718	< 0.001	-26.718	> 0.100	0.113
USD/CHF	< 0.001	-11.007	< 0.001	-22.894	> 0.100	0.041
USD/GBP	< 0.001	-9.439	< 0.001	-24.724	> 0.100	0.220
USD/JPY	< 0.001	-9.351	< 0.001	-25.875	> 0.100	0.081
USD/SEK	< 0.001	-10.126	< 0.001	-28.309	> 0.100	0.135
USD/ZAR	< 0.001	-10.091	< 0.001	-26.948	> 0.100	0.151
USD/MXN	< 0.001	-10.104	< 0.001	-26.254	> 0.100	0.105
USD/XBT	< 0.001	-8.300	< 0.001	-29.383	> 0.100	0.172

The above table shows test statistics and associated p-values from performing stationarity tests of logarithms of the spot exchange rate pairs used in this analysis, during the sample time period of 1 May 2013 – 31 December 2015. Reported are results from the augmented Dickey-Fuller (ADF), Phillips-Perron (PP), and Kwiatkowski et al. (1992) (KPSS) tests. ADF and PP test for the null of a unit root. The KPSS tests the null of stationarity. Results are presented for log exchange rates in levels (Panel A) and in first differences (Panel B). Note that "$>$" and "$<$" mean that the test statistics are above or below the maximum and minimum tabulated critical values, resp. The minimum tabulated critical values are at the 1% confidence level for the KPSS test and and at 0.1% for the ADF and PP tests. The maximum tabulated critical value for the KPSS test is at 10%.

Table 7: Precision of Bitcoin (XBT) Exchange Rate Observations

	(1) % At	(2) #Obs(Mean)	(3) #Obs(Median)	(4) #Out(Mean)	(5) #Out(Med)	(6) #Out(Mode)	(7) Earliest	(8) Median	(9) Latest
USD/XBT	88.11%	64.90	35	0.24	0	0	-0.02	0.00	0.02
EUR/XBT	65.23%	12.93	6	0.80	0	0	-0.01	0.00	0.02
AUD/XBT	10.79%	3.18	3	2.41	3	3	8.65	10.94	11.57
CAD/XBT	6.59%	3.09	3	2.55	3	3	10.83	11.45	11.77
CHF/XBT	2.20%	3.04	3	2.85	3	3	-0.97	1.00	4.66
GBP/XBT	24.58%	3.47	3	1.75	2	3	0.01	11.05	11.72
JPY/XBT	13.39%	3.47	3	2.31	3	3	1.33	7.49	10.08
SEK/XBT	1.10%	3.02	3	2.83	3	3	7.43	9.27	10.63
ZAR/XBT	1.30%	3.02	3	2.73	3	3	4.57	6.42	7.96
MXN/XBT	0.00%	3.00	3	3.00	3	3	9.10	10.11	11.05

This table shows summary statistics regarding the precision of our estimates of noon Bitcoin (XBT) exchange rates. The goal is to capture the noon exchange rate as a volume-weighted average of transaction prices occurring between 11:59am and 12:01pm ET. To limit the effects of extreme prices, a minimum of three transaction prices is required to calculate the noon exchange rate for each day. If less than three transactions are found to occur between 11:59am and 12:01pm ET, our algorithm then takes the next-closest observations to the target window in terms of time, until a minimum of three observations are found. Column 1 in the table shows the percentage of daily noon rates for which all observations used in its calculation are observed within the target time window. Columns 2-3 show the mean and median number of observations used to calculate each noon rate. Columns 4-6 show the mean, median, and mode number of observations outside the target window used to calculate the noon rate. Columns 7-9 show the median of three variables: the furthest-away time before the target window, the median time away from the target window, and the further-away time after the target window.