The Role of Binance in Bitcoin Volatility Transmission

Carol Alexander*, Daniel F. Heck† and Andreas Kaeck‡

July 2, 2021

Abstract

We analyse high-frequency realised volatility dynamics and spillovers in the bitcoin market, focusing on two pairs: bitcoin against the US dollar (the main fiat-crypto pair) and trading bitcoin against tether (the main crypto-crypto pair). We find that the tether-margined perpetual contract on Binance is clearly the main source of volatility, continuously transmitting strong flows to all other instruments and receiving only a little volatility. Moreover, we find that (i) during US trading hours, traders pay more attention and are more reactive to prevailing market conditions when updating their expectations and (ii) the crypto market exhibits a higher interconnectedness when traditional Western stock markets are open. Our results highlight that regulators should not only consider spot exchanges offering bitcoin-fiat trading but also the tether-margined derivatives products available on most unregulated exchanges, most importantly Binance.

Keywords: Bitcoin ETF; Exchange-Traded Funds; Realised Volatility; Volatility Transmission.

JEL classification: C22, C5, E42, F31, G1, G2

*University of Sussex Business School. Email: c.alexander@sussex.ac.uk
†University of Sussex Business School. Email: d.heck@sussex.ac.uk
‡University of Sussex Business School. Email: a.kaeck@sussex.ac.uk

1 Introduction

During the most recent bull run in the first half of 2021, several of the largest financial institutions have started entering the crypto market, launching digital assets groups.[1] At the same time, the Chicago Mercantile Exchange (CME) expanded their product spectrum by ether futures and options as well as micro bitcoin futures which most recently hit the milestone of one million contracts traded – within just two months. Unsurprisingly, there have also been multiple bitcoin ETF filings, most notably by Fidelity.[2] So far however, none of the applications has been approved by the Securities and Exchange Commission (SEC), mainly due to inadequate investor protection and insufficient resistance to fraud and manipulation.

Despite continuous growth, the regulated crypto market remains tiny compared to the unregulated exchanges, where especially the crypto derivatives – perpetual swaps and futures – are traded very actively. All these exchanges offer two types of derivatives products, namely inverse (so called coin-margined) and linear (USDT-margined) contracts. The former are rather complicated products: for ease of use, they are quoted in USD but their actual quote currency is bitcoin so they are contracts on USD/BTC instead of BTC/USD. This way, unregulated exchanges avoid dealing with fiat currencies and their onboarding process. The linear products on the other hand are exactly like derivatives in traditional asset classes, however they are not settled and traded in traditional fiat currencies, but using tether. This is the largest stablecoin – at the time of writing, it was the third largest cryptocurrency – whose value is pegged to the US dollar. Its issuer, Tether Limited, originally claimed to hold one US dollar for each token of tether. In May 2021 however, after an investigation by the New York Attorney General, they reported that only 2.9% of all tokens are backed by cash reserves, but they are still unwilling to undergo a reliable third-party audit.[3] Moreover, there have been repeated allegations against tether that it has been used to inflate the prices of cryptocurrencies while being unbacked (Griffin and Shams, 2020). Nevertheless, tether's usage in the crypto market has continuously been growing – mainly as quote currency for other cryptocurrencies and margin unit for (potentially highly-leveraged) derivatives products. Therefore, the stablecoin definitely requires more attention, both from academics and practicioners as well as regulators.

We try to shed more light on tether's ability for contagion and volatility transmission within the bitcoin market. To this end, we use high-frequency data to analyse realised volatility spillovers among major bitcoin exchanges, including both USD and USDT spot pairs as well as inverse and linear perpetual contracts. Using a multivariate version of the Logarithmic Multiplicative Error Model which is particularly well suited for modelling high-frequency financial time series, we find that the tether-margined perpetual contract on Binance is clearly the main emitter of volatility. Throughout the day, it continuously transmits strong flows to all other instruments. However, the strength of these flows varies over the course of the day. They are lowest during Asian trading, then intensify substantially in European trading and finally reach their maximum during US trading hours. Out of all instruments included in our analysis, the tether-based contract also receives the lowest volatility flows, which apart

[1] See Coindesk and Bloomberg.
[2] See CNBC.
[3] See Financial Times.

from one minor exception holds throughout the day. The remaining crypto instruments generate much weaker volatility flows and therefore have a significantly lower contagion potential. Moreover, we find that the overall volatility flows as well as the total short-term persistence strengthen over the course of the day. This indicates that (i) during US trading hours, traders pay more attention and are more reactive to prevailing market conditions when updating their expectations and (ii) the crypto market exhibits a higher interconnectedness when traditional Western stock markets are open. Finally, as a by-product to our spillover analysis, we document an intraday pattern of trading volume and realised volatility that seems to be similar to FX markets as documented by Andersen and Bollerslev (1997). On all instruments, we also identify distinct volume and volatility spikes every four hours – most notably at midnight and 16:00 UTC – which we attribute to the funding of perpetuals.

These volatility transmission results on their own might not be particularly alarming. However, in light of the recurring suspicions of misconduct against tether and more recently also Binance – in the US, they are under investigation by the Justice Department, while the German watchdog warned Binance and both the UK and Japan started cracking down on the exchange – and the still very loose regulation of the whole crypto market, our results become rather concerning.[4] In particular, they highlight that academics and regulators should expand their focus from bitcoin-fiat pairs to bitcoin-tether products. When deciding on the numerous bitcoin ETF applications, the SEC should not only consider spot exchanges offering bitcoin-fiat trading – some of whom voluntarily comply with financial regulation, such as Coinbase or Gemini – but also the tether-margined derivatives products available on most unregulated exchanges, especially Binance who do not convey the impression of being particularly willing to cooperate with financial authorities.[5]

Our study makes two major contributions to the existing literature: First, we base our analysis of volatility flows and co-movement on high-frequency data. In particular, we use second-by-second observations to compute realised volatiliy at the five-minute frequency, which due to microstructure noise is the highest frequency allowing a reliable transmission analysis. All other studies on volatility spillovers among different cryptocurrencies (Yi et al., 2018; Katsiampa et al., 2019; Wang and Ngene, 2020; Caporale et al., 2021; Sensoy et al., 2021) or between individual crypto-exchanges (Ji et al., 2021) rely on daily observations and/or focus exclusively on the much smaller spot market. However, we show both that volatility flows among crypto instruments exhibit significant intraday-variation and that the tether-based perpetual contract on Binance is the main source of volatility. Therefore, analysing volatility spillovers at a daily frequency or omitting the bitcoin derivatives market can lead to erroneous conclusions.

Moreover, we highlight the important role tether plays within the crypto space. Besides a few recent studies such as Griffin and Shams (2020); Ji et al. (2021); Baur and Hoang (2021), the stablecoin has not attracted significant attention from academics or regulators, despite its huge and continuously growing market capitalization. However, based on high-frequency data, we find that trading bitcoin against tether is the main source of volatility. In particular, the perpetual contract that allows trading the two cryptocurrencies with a very high leverage generates and transmits a large amount of volatility

[4] See Reuters on 13 May, Reuters on 29 April and Yahoo Finance on 28 June.

[5] See Coindesk on 8 May 2020.

to all other instruments. The BTC/USD pair on the other hand (both spot and perpetual) emits much less volatility and mainly receives flows from the tether-based perpetual swap.

The remainder of this paper is organized as follows: Section 2 describes the exchanges and instruments included in our analysis and their second-level price and volume data; Section 3 explains the econometric framework; Section 4 describes the realised volatility calculation and examines its intraday pattern across instruments; Section 5 presents our results on volatility transmission, both within and across different instrument groups, and their intraday variation; Section 6 summarises and concludes.

2 Exchanges, Instruments and Data

Here we describe the characteristics and trading volumes on all instruments included in our analysis. We adopt the standard ticker BTC for bitcoin.We only admit the major (well-established) bitcoin spot exchanges (Bitstamp, Coinbase, Huobi, Kraken) as well as major (unregulated) crypto derivatives exchanges into our analysis (Binance, Bybit). However, we exclude futures from our analysis and focus on the very popular perpetual swaps (or perpetual for short) which generally exhibit much higher trading volumes. These contracts are currency swaps between two currencies (either crypto–crypto or crypto–fiat) where the only cash-flows are between perpetual fixed and floating legs and these are usually exchanged three times a day. They combine the features of both futures and spot positions in that they do not expire before being closed out (similar to a spot position) and they allow very high-leverage trading (like a futures contract). In contrast to ordinary futures, perpetual contracts are not exposed to any roll-over risk and the basis is very much smaller than it is for futures.

The major focus of our analysis lies on two currency pairs: BTC/USD and BTC/USDT. The ticker USDT refers to tether, a so-called stablecoin whose value is pegged to the US dollar. Tether is controversial since its issuer falsely claimed to hold one US dollar physically for each token of tether.[6] Moreover, the stablecoin was allegedly used to inflate the price of bitcoin artifically (Griffin and Shams, 2020). Nevertheless, tether is still widely used – mainly as quote currency for both crypto spot pairs and margin and settlement currency for derivatives instruments – and has continuously been among the largest cryptocurrencies over the last years. Recently, its market capitalization which for a stablecoin is equal to the number of tokens in circulation even exceeded \$60 bn.[7]

Since the crypto market evolves at an extremely fast pace, we are interested in high-frequency inter-exchange volatility flows, rather than any long-term dynamics. Therefore, we choose a one-second data sampling frequency so that we are able to compute realised volatility every five minutes. We focus on the recent bull market from 1 January to 31 March 2021, where the price of bitcoin rose by almost 100% from \$30,000 to around \$60,000. Figure 1 shows the detailed price evolution over this three-months period.

Table 1 shows the product specifications for the different perpetual swaps. The USD-contracts are of inverse type, i.e. their base currency is BTC so they are contracts on USD/BTC instead of BTC/USD. The tether-contracts however are ordinary linear products, that is they are settled in

[6] See this release of the New York Attorney General.
[7] See Yahoo finance.

Figure 1 Bitcoin Price from 1 January to 31 March 2021

Note: The figure shows the bitcoin price on Coinbase in USD, from 1 January to 31 March 2021.

tether. As is usual with crypto products, the perpetual swaps allow very high leverage and can be traded 24/7.

A key point to note from Table 1 is that the fees on all the unregulated derivatives exchanges follow (at least partly) a maker-taker model. That is, orders that add liquidity to the book (non-marketable limit orders) are charged less than orders that reduce liquidity (market orders or marketable limit orders). In the case of Bybit, maker fees are even negative, i.e. liquidity-increasing orders are refunded a certain percentage of the order size, independent of the user's past trading volume. For their USD-perpetual, Binance follow a tiered maker-taker model refusing ordinary users maker rebates. Only VIP investors with a 30-day trading volume of at least 50,000 XBT and a balance of 2,000 of their in-house token (called Binance Coin; BNB) obtain maker rebates. But even in their respective top tier – trading volume exceeding 750,000 XBT and a balance of more than 11,000 BNB – the rebates are 0.9 basis points and thus, significantly smaller than on Bybit. For the linear USDT-perpetual, Binance does not offer maker rebates at all, VIP investors can only achieve zero maker fees. Similarly, the spot exchanges do not offer maker rebates and follow a volume-dependent fees schedule of at most 50bps (Coinbase and Bitstamp), 26bps (Kraken), 20bps (Huobi) and 10bps (Binance). Bitstamp is the only spot exchange not offering discounts for liquidity-increasing orders.

All three of these bitcoin derivatives exchanges are still completely unregulated. There is no supervisory authority establishing any rules to prevent malpractice, misconduct and manipulation.[8] Admittedly, it is a very challenging task to regulate these venues. If the Commodity Futures Trading Commission (CFTC) went after some exchange and forced them to implement proper KYC and AML-mechanisms, the majority of the exchange's users would simply switch to some other venue still offering very loose or no regulation. This already happened when the CFTC charged BitMEX – one of the largest crypto derivatives exchanges at that time – for illegaly operating a cryptocurrency derivatives

[8]In fact, some exchanges such as Huobi are even unauthorized and explicitly banned in the US. In practice however, this ban can be circumvented quite easily using a virtual private network (VPN).

Table 1 Perpetuals Specifications

	USD Contracts		USDT Contracts
	Binance	Bybit	Binance
Type	Inverse	Inverse	Linear
Contract Size	100 USD	1 USD	0.001 BTC
Margin Requirement	0.8%*	1%	0.8%*
Settlement Currency	BTC	BTC	USDT
Trading Days	24/7	24/7	24/7
Funding Frequency	8 hrs	8 hrs	8 hrs
Fees (maker/taker)	1/5	-2.5/7.5	2/4
Tick Size	0.1 USD	0.5 USD	0.01 USDT

Note: The table shows the main specifications of the perpetual contracts included in our analysis. All fees reported here are in basis points. *The leverage on Binance depends on the notional value of the position. The larger the position, the lower the leverage allowed.

trading platform and AML violations.[9] However, even though we cannot rule out the possibility of price or volume manipulation, this does not prevent a proper spillover analysis. If anything, the lack of supervision makes our study more interesting for market participants and regulators.

We retrieved data on all instruments from coinAPI.[10] The dataset consists of second-by-second price and volume data and covers the period from 1 January to 31 March 2021, summing up to more than 7.5 million observations per instrument and almost 60 million in total. Table 2 provides summary statistics on the trading activity of the different instruments included in our analysis. We can see that trading on perpetuals (both USD and USDT) and BTC/USDT spot pairs is highly active. Transactions occur in at least three out of four seconds – for the tether-instruments, this fraction is even more than 95% – and the average daily volume (ADV) exceeds $1.5 bn. The highest ADV occurs on the Binance USDT-perpetual with even more than $15 bn. Out of the three USD spot pairs, only Coinbase is able to keep up. Transactions are conducted in 84% of one-second intervals and its ADV is about $1.2 bn. With trading volumes of less than half a billion US dollar and transactions occurring only in about every fourth second, Bitstamp and Kraken lag far behind.

It is well documented that trading volume exhibits a certain intraday pattern. In North American equity markets, many studies find a U-shaped pattern (Jain and Joh, 1988; McInish and Wood, 1990), while Cai et al. (2004) and Ozenbas (2008) document a more M-shaped behaviour on the London Stock Exchange. For major currency pairs, Danielsson and Payne (2001) and McGroarty et al. (2009) find a similar M-shaped volume pattern with peaks at London and New York opening times. Figure 2 shows the intraday pattern of trading volume for BTC/USD and BTC/USDT, both on spot exchanges and perpetual contracts, measured as the median five-minute trading volume over the period from 1 January to 31 March 2021.[11] As can be seen, volume follows a similar pattern on spot and perpetuals. First, it evolves in a U-shape with very distinct peaks at midnight and 16:00 UTC. The most extreme spike can be observed on the Huobi USDT spot pair, where the median trading volume increases

[9] See this CFTC release on 1 October 2020.

[10] Data and software provider coinAPI links with hundreds of crypto spot and derivatives exchanges, offering historical and streaming order-book and trades in tick-by-tick or highest granularity data from all major centralized and decentralized exchanges.

[11] Note that the number of transactions shows an analogous intraday pattern.

Table 2 Data Statistics

Type	Currency	Instrument	Count	ADV	MaxDV
Spot	USD	Bitstamp	1,867,238 (24%)	392.92	1,364.60
		Coinbase	6,554,500 (84%)	1,187.42	3,499.38
		Kraken	1,777,536 (23%)	345.03	893.02
	USDT	Binance	7,716,626 (99%)	3,858.29	8,410.34
		Huobi	7,420,413 (95%)	1,586.62	4,194.10
Perpetuals	USD	Binance	5,798,339 (75%)	5,481.04	15,084.64
		Bybit	6,748,841 (87%)	7,368.11	17,755.14
	USDT	Binance	7,712,699 (99%)	15,006.25	36,297.44

Note: The table shows the number of second intervals where at least one trade was conducted (Count, in absolute terms and as percentage of the total number of second-intervals), the average daily volume (ADV, in million USD) and the maximum daily volume (MaxDV, in million USD), during the period from 1 January to 31 March 2021.

fivefold from around \$5 m to almost \$25 m at 16:00 UTC. After the afternoon spike, trading volume on all instruments continuously decreases to an average level. The three USD spot pairs reach their minimum between 09:00 and 10:00 UTC, while the volumes on perpetuals and USDT spot pairs seem to be at their lowest in the early morning around 03:00 UTC. Consequently, we can only partly confirm the results of Eross et al. (2019) who document an inverted U-shape on Bitstamp with a peak at around 14:00 UTC. However, they consider the period from 2014 to 2017 and therefore, the difference in intraday trading volume might be explained by a significant evolution of the crypto market over the last three years.

The two graphs of figure 2 show another interesting feature. Trading volume exhibits distinct spikes in the first five minutes of every fourth hour, e.g. from 08:00 to 08:05 UTC. These peaks are especially pronounced for the perpetual swaps. Moreover, it seems as if these products also show increased trading volume at the beginning of each hour, albeit to a lesser extent. There is only one explanation for these spikes that we can think of, namely the funding of perpetuals. While the contracts considered here (Binance, Bybit) are funded at 00:00, 08:00 and 16:00 UTC, other major perpetuals such as the BitMEX one exchange cash flows at 04:00, 12:00 and 20:00 UTC. Possibly, some market participants modify their positions across multiple exchanges just after funding on some perpetual has occurred or they take advantage of some mispricing, either between different perpetuals or between spot and perpetuals. Using historical data on the Binance funding rates and basic linear regressions, we do however not find evidence that the height of these volume spikes is related to the funding rate.

3 Methodology: Multiplicative Error Model

Since volatility is a non-negative variable, we base our analysis on the Multiplicative Error Model (MEM) introduced in Engle and Russell (1998).[12] This model implicitly guarantees non-negativity of the variables of interest and overcomes certain difficulties and inefficiencies of the standard Gaussian

[12]Originally, the MEM has been proposed to model durations between transactions and it is therefore often called Autoregressive Conditional Duration (ACD) model. However, throughout this paper, we refer to the model and its extensions by the more general term Multiplicative Error Model.

Figure 2 Volume Intraday Pattern

Note: The figure shows the intraday pattern of five-minute trading volume (in million USD) for USD spot pairs (upper graph), USDT spot pairs (middle graph) and perpetuals (lower graph), measured as the median five-minute volume over the period from 1 January to 31 March 2021. All times are in UTC.

approach when modelling non-negative time series (Engle, 2002). Formally, the MEM decomposes the variable of interest into a product of its conditional mean and an error term which has unit mean and follows a distribution with non-negative support. The conditional mean is usually assumed to be autoregressive, depending both on its own lagged values and those of the variable of interest. Since we analyse realised volatility, we also include an asymmetric response component to capture the leverage effect typically observed in financial markets (Bollerslev et al., 2006).

In the original MEM however, certain parameter restrictions have to be imposed to ensure non-negativity which exacerbates model estimation and interpretation. Therefore, we choose to use the LogMEM$_1$ introduced by Bauwens and Giot (2000), where the conditional mean is replaced by its logarithm, making any non-negativity constraints obsolete. Our model is therefore given by

$$\begin{aligned} x_t &= \mu_t \, \varepsilon_t \\ \log \mu_t &= \omega + \sum_{j=1}^{p} \alpha_j \log x_{t-j} + \gamma \log x_{t-1}^{-} + \sum_{j=1}^{q} \beta_j \log \mu_{t-j} \end{aligned} \tag{1}$$

where $\log x_{t-1}^{-} = \log x_{t-1}$ if the return of the respective interval is negative and zero otherwise. As required, exogenous variables such as dummies to capture time variation may also be included in the conditional mean equation.

Finally, we need to specify the conditional distribution of the innovations. As shown by Allen et al. (2014), the log-normal distribution yields consistency and asymptotic normality of the Quasi-Maximum Likelihood (QML) estimator in the LogMEM$_1$ and exhibits superior finite sample properties compared to other widely-used error distributions. However, since we analyse high-frequency time series, it is quite likely to encounter zero observations, which the log-normal distribution is not able to capture. Therefore, we follow Nguyen et al. (2020) and apply the zero-augmented version of the log-normal distribution as proposed by Hautsch et al. (2014). Its density function is given by

$$f(x) = (1 - p^{+})\delta(x) + p^{+} \frac{1}{\sqrt{2\pi}\, x\, s} \exp\left(-\frac{1}{2} \left(\frac{\log(x) - m}{s} \right)^2 \right) \mathbb{1}_{\{x>0\}}$$

where $\delta(x)$ is a point probability mass at zero, $m = -0.5s^2 - \log p^{+}$ due to the requirement of unit mean and p^{+} denotes the empirical ratio of zero observations in the time series. However, since the logarithm is only defined for strictly positive values, we have to modify the conditional mean specification in (1) by adding auxiliary parameters α^0 which capture the effect of zero volatilites on the conditional mean, i.e.

$$\log \mu_t = \omega + \sum_{j=1}^{p} \alpha_j \log x_{t-j} \, \mathbb{1}_{\{x_{t-j}>0\}} + \sum_{j=1}^{p} \alpha_j^0 \, \mathbb{1}_{\{x_{t-j}=0\}} + \gamma \log x_{t-1}^{-} + \sum_{j=1}^{q} \beta_j \log \mu_{t-j} \tag{2}$$

This specification can easily be extended to multiple dimensions to obtain the vector Logarithmic

Multiplicative Error Model (vLogMEM$_1$). It is given by

$$
\begin{aligned}
\boldsymbol{x}_t &= \boldsymbol{\mu}_t \odot \boldsymbol{\varepsilon}_t \\
\log \boldsymbol{\mu}_t &= \mathbf{w} + \sum_{j=1}^{p} \mathbf{A}_j \left(\log \boldsymbol{x}_{t-j} \odot \mathbb{1}_{\{\boldsymbol{x}_{t-j}>0\}}\right) + \sum_{j=1}^{p} \mathbf{A}_j^0 \mathbb{1}_{\{\boldsymbol{x}_{t-j}=0\}} + \boldsymbol{\Gamma} \log \boldsymbol{x}_{t-1}^{-} + \sum_{j=1}^{q} \mathbf{B}_j \log \boldsymbol{\mu}_{t-j}
\end{aligned} \tag{3}
$$

where $\odot$ denotes the Hadamard (element-by-element) product and the indicator functions should be interpreted component-wise. However, the specification of the error distribution is much more challenging in the multivariate model, since the variables of interest are likely to be interdependent. There are mainly three ways of dealing with this. First, we could choose the multivariate log-normal distribution, which is one of the very few multivariate distributions with non-negative support (Taylor and Xu, 2017). However, this distribution is not able to capture probability mass at zero and an appropriate adaption as in the univariate case is far from trivial. Second, we could apply copulae to link the univariate distributions of variables (Cipollini et al., 2017). Due to the non-trivial fraction of valid zero observations however, this approach is quite challenging, since we would have to estimate the zero-augmented mixture distributions and the copula jointly (Nguyen et al., 2020). Therefore, we choose to follow the third option where the innovations of the individual variables are assumed to be orthogonal and only lagged interdependence is allowed. This assumption requires both the innovations' covariance matrix and the long-term persistence matrix $\mathbf{B}$ to be diagonal, which makes this approach equivalent to fitting the above univariate models for each variable individually, with the lagged values of the remaining instruments as exogenous variables. Even though this equation-by-equation approach leads to a loss in efficiency, it still yields consistent estimates in a QML context. Moreover, compared to a model allowing full interdependence, the procedure limits the number of parameters to be estimated and thus sidesteps the curse of dimensionality (Escribano and Sucarrat, 2018).

4 Realised Volatility

Many studies have already investigated volatility spillovers among different cryptocurrencies or individual crypto-exchanges (Katsiampa et al., 2019; Cheah et al., 2018). However, most of these studies rely on daily observations and consider only the much smaller spot market. To the best of our knowledge, no study has yet analysed high-frequency volatility flows within the crypto spot and derivatives market and we aim to fill this gap with this study. Therefore, using second-by-second observations, we calculate realised volatility at a five-minute frequency which seems to be a reasonable trade-off between too much microstructure noise at a higher frequency and potentially missing significant volatility flows at lower frequencies (Andersen, 2000; Nguyen et al., 2020).

Since there is still a lot of noise in the data, we calculate realised volatility using the robust estimator based on pre-averaging (Jacod et al., 2009). Here, the one-second log returns are locally smoothed using a weighted average, which reduces the influence of microstructure noise to some extent. In particular, the pre-averaged returns are given by

$$\bar{r}_t = \sum_{j=1}^{k_n} g\left(\frac{j}{k_n}\right) r_{t+j}$$

where r_t denotes the ordinary (not pre-averaged) log returns and g is a real-valued weighting function, commonly chosen as $g(x) = \min(x, 1-x)$.[13] Following the empirical analysis of Hautsch and Podolskij (2013), the bandwidth k_n of the local pre-averaging window is chosen as $\lceil \theta\sqrt{n} \rceil$ where n is the number of observations (in our case, $n = 300$) and $\theta \in [0.3, 0.6]$. The exact choice of the smoothing parameter θ mainly depends on the sampling frequency and the liquidity of the asset. Since we are interested in multivariate spillovers, rather than individual volatility estimates that are as accurate as possible, we set θ consistently as 0.4 for all instruments. Therefore, the optimal bandwidth is $\lceil 0.4\sqrt{300} \rceil = 7$.

For each five-minute interval, the pre-averaging estimator of realised variance is then calculated as the sum of squared pre-averaged one-second returns. Finally, to obtain the pre-averaged realised volatility, we take the square root of the realised variance and annualise it by a factor of $\sqrt{12 \times 24 \times 365}$.

If no transaction has been executed within a one-second interval, we fill this gap with the previous price which leads to a zero return and therefore induces a downward bias in realised volatility. However, it seems intuitive that an instrument with lower trading activity should also exhibit a lower realised volatility. In addition, the pre-averaging reduces the influence of such zero returns on the realised volatility measure. Therefore, using the last available price is, in our opinion, the most appropriate way to deal with one-second intervals in which no transaction has taken place.

Since the crypto market is still not as liquid as other more established asset classes, it is possible that for some five-minute intervals, only a very limited number of transactions has been carried out. Therefore, in order to obtain reliable volatility estimates, we apply a threshold on trading activity: if there are only transactions in less than 20% of the one-second intervals, we set the volatility estimate of that five-minute period equal to zero. Following table 2, this threshold will not be of too much concern for Coinbase, Huobi and the three Binance instruments, but it may be relevant for Bitstamp and Kraken. However, since we rely on the zero-augmented log-normal distribution in the LogMEM$_1$ and explicitly account for zero observations in the conditional mean specification (2), it will not be a major issue if these two exchanges exhibit a comparatively large number of zero values.

Table 3 shows summary statistics on the five-minute realised volatility for the period from 1 January to 31 March 2021. We can see that the average levels of volatility on Coinbase, Bitstamp, Huobi, Bybit and the three Binance instruments are not too different, while that of Kraken is significantly lower, probably due to comparatively low trading activity. For the same reason, Bitstamp and Kraken show rather a high standard deviation of more than 60% – compared to about 43% on the six remaining instruments – paired with a zero 25%-quantile. We also note that all eight instruments show a minimum volatility of zero, implying that our trading activity threshold is not met at least once for all instruments. Overall however, except for Bitstamp and Kraken, the realised volatility statistics

[13] The only requirements on the weighting function $g : [0,1] \to \mathbb{R}$ are that it is continuous, piecewise continuously differentiable with piecewise Lipschitz derivative and $g(0) = g(1) = 0$. Since the tent-shaped function $g(x) = \min(x, 1-x)$ is the simplest functions meeting these criteria, it is commonly chosen as weighting function.

look quite similar. Finally, it is striking that Kraken shows an extremely high maximum volatility of 63. This value is attained on 22 February 2021 between 14:20 and 14:25 UTC and therefore, figure 3 displays the prices on Kraken and the Binance USDT-perpetual at this time. As we can see, the price on Kraken jumps up and down with an increasing amplitude – between 14:24 and 14:25, the jump size is even $4,000 – resulting in a very high realised volatility. The price on Binance on the other hand evolves very smoothly without any visible jumps. Therefore, we suspect that this extremely large bid-ask bounce might be caused by some exchange-specific issues, like server problems for example. To reduce, but not eliminate, the impact of such anomalous outliers, we will winsorize the data before our empirical analysis.

Table 3 Realised Volatility Statistics

	Coinbase	Bitstamp	Kraken	Binance^S	Huobi	Binance^T	Bybit	Binance$^\$$
Mean	0.6985	0.6217	0.4709	0.7237	0.7299	0.7630	0.6146	0.7068
Std	0.4255	0.6102	0.6837	0.4155	0.4147	0.4371	0.4408	0.4409
Min	0	0	0	0	0	0	0	0
25%	0.4344	0	0	0.4650	0.4721	0.4909	0.3347	0.4311
50%	0.5945	0.6262	0.4012	0.6235	0.6302	0.6534	0.5044	0.5990
75%	0.8326	0.9504	0.7358	0.8541	0.8620	0.8972	0.7562	0.8464
Max	11.6779	8.0250	63.3657	8.5440	7.8419	7.4741	6.6263	7.7105

Note: The table shows summary statistics on the five-minute realised volatility over the period from 1 January to 31 March 2021. Binance^S, Binance^T and Binance$^\$$ denote the BTC/USD spot pair, the USDT-perpetual and the USD-perpetual, respectively, on Binance.

Figure 3 Bitcoin Price on 22 February 2021

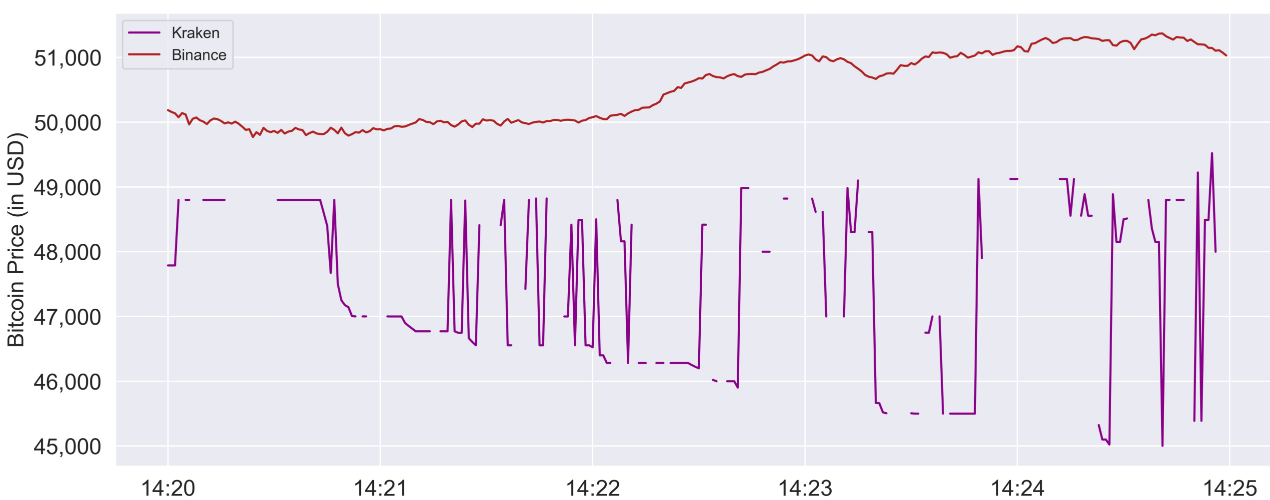

Note: The figure shows the bitcoin price in USD on Kraken (magenta line) and Binance (red line) on 22 February 2021 from 14:20 to 14:25 UTC.

Closely related to trading volume, realised volatility also exhibits a certain intraday pattern (Andersen and Bollerslev, 1997). As can be seen in figure 4, five-minute realised volatility shows a pattern very similar to that of trading volume. For both spot pairs and perpetuals, it first follows a U-shape with distinct peaks at midnight and 16:00 UTC and smaller spikes at 04:00 and 08:00 UTC. After 16:00, it slowly declines to between 50% and 80%, with again a visible spike at 20:00 UTC. In the upper graph, we can also see that Bitstamp and Kraken exhibit a lower volatility than Coinbase and the remaining

USDT spot pairs and perpetual contracts. The difference is especially pronounced from midnight until 11:00 UTC, i.e. during Asian and early European trading hours. Therefore, we suppose that mainly American and European traders are active on these two exchanges, leading to a reduced activity and consequently to a lower volatility – since we are filling intervals where no trade occurred with the last available price – during the times American and European equity markets are usually closed. Overall however, the eight instruments exhibit a quite similar volatility intraday pattern. Finally, we note that the extremely large volatility on Kraken around 14:20 UTC is caused by the exchange-specific outlier discussed above and visualised in figure 3.

In fact, the intraday pattern of volatility seems similar to the results of Andersen and Bollerslev (1997) who examine absolute returns for the Deutsche Mark/USD pair. They find that volatility starts at a relatively high level at midnight, then slowly decays and reaches its minimum around 03:00 UTC. After that, trading activity increases to a local maximum around opening of European markets. The overall maximum however is attained at the opening of US markets around 13:00 UTC. Afterwards, the volatility slowly declines and only starts to pick up again around 21:00 UTC.

It is important to account for this deterministic intraday pattern, otherwise the results from the MEM could be skewed. We do so in a non-parametric way by dividing each realised volatility observation by the average realised volatility of the respective five-minute interval, documented in figure 4.[14] Alternatively, we could also model the intraday pattern explicitly in the MEM, but this would increase the number of parameters to be estimated even more. To account for time-variation in the modelled variables, Nguyen et al. (2020) use the 250-day moving average as diurnal adjustment factor. Due to data constraints however, this is not possible in our study and we simply use the average over the complete period from 1 January to 31 March 2021. Since our sample period is rather short compared to that of Nguyen et al. (2020), time-variation is not of too much concern.

As mentioned above, we winsorize the top 0.05% of our diurnally-adjusted realised volatility values, i.e. all observations greater than the 99.95%-quantile are set equal to this quantile. Figure 5 shows the distribution of the five-minute realised volatility across the eight instruments, both before and after diurnal adjustment and winsorizing. Due to the quite large number of zero observations on Kraken and Bitstamp, we excluded these zero values from the histograms, allowing a better comparison across instruments. The figure confirms the summary statistics and graphs from above: the realised volatilities on Coinbase, Huobi, Bybit and the three Binance instruments exhibit quite similar, rather narrow shapes with a comparatively low variation. Only Bybit seems to differ slightly – its empirical density attains a maximum of around 0.9, compared to between 1.1 and 1.2 on Coinbase, Huobi and Binance. Due to rather low trading activity, Bitstamp and Kraken exhibit quite a high portion of zero values and consequently, their empirical densities are continuously less than those of the remaining six instruments. Overall however, the histograms seem appropriate and we can continue with our empirical analysis.

[14] To reduce the impact of extreme outliers, we could use the median five-minute volatility as diurnal adjustment factor. However, Kraken exhibits zero median values in the early UTC morning and therefore, we rely on the average five-minute volatility. Apart from Kraken in the early morning, average and median intraday patterns have a very similar shape anyway.

Figure 4 Volatility Intraday Pattern

Note: The figure shows the intraday pattern of 5-minute realised volatility (in million USD) for USD spot pairs (upper graph) as well as USDT spot pairs and perpetuals (lower graph), measured as the average five-minute realised volatility over the period from 1 January to 31 March 2021. All times are in UTC.

5 Volatility Transmission Results

5.1 Univariate Dynamics

To detect any differences among instruments, we first fit the univariate LogMEM$_1$ given in equations (1) and (2) to the diurnally-adjusted five-minute realised volatility of each instrument. We choose to include only one lag ($p = q = 1$). As argued in detail by Nguyen et al. (2020), the maximum number of lags is expected to be three, but in practice it will be lower due to finite sample issues for example.[15]

Table 4 reports the resulting parameter estimates, together with the log-likelihood value, the

[15] As a robustness check, we have conducted a thorough analysis on the number of lags included in the LogMEM$_1$. In general, the total persistence – the sum of all α and β – as well as the distribution parameter s are independent of p and q. The mean squared error and the Ljung-Box test statistic change only very marginally, while log-likelihood and BIC improve only when increasing the number of lags from one to two. The main difference seems to be the "spread" of the persistence across the α and β parameters, which is in line with Nguyen et al. (2020).

Figure 5 Histograms of Five-Minute Realised Volatility

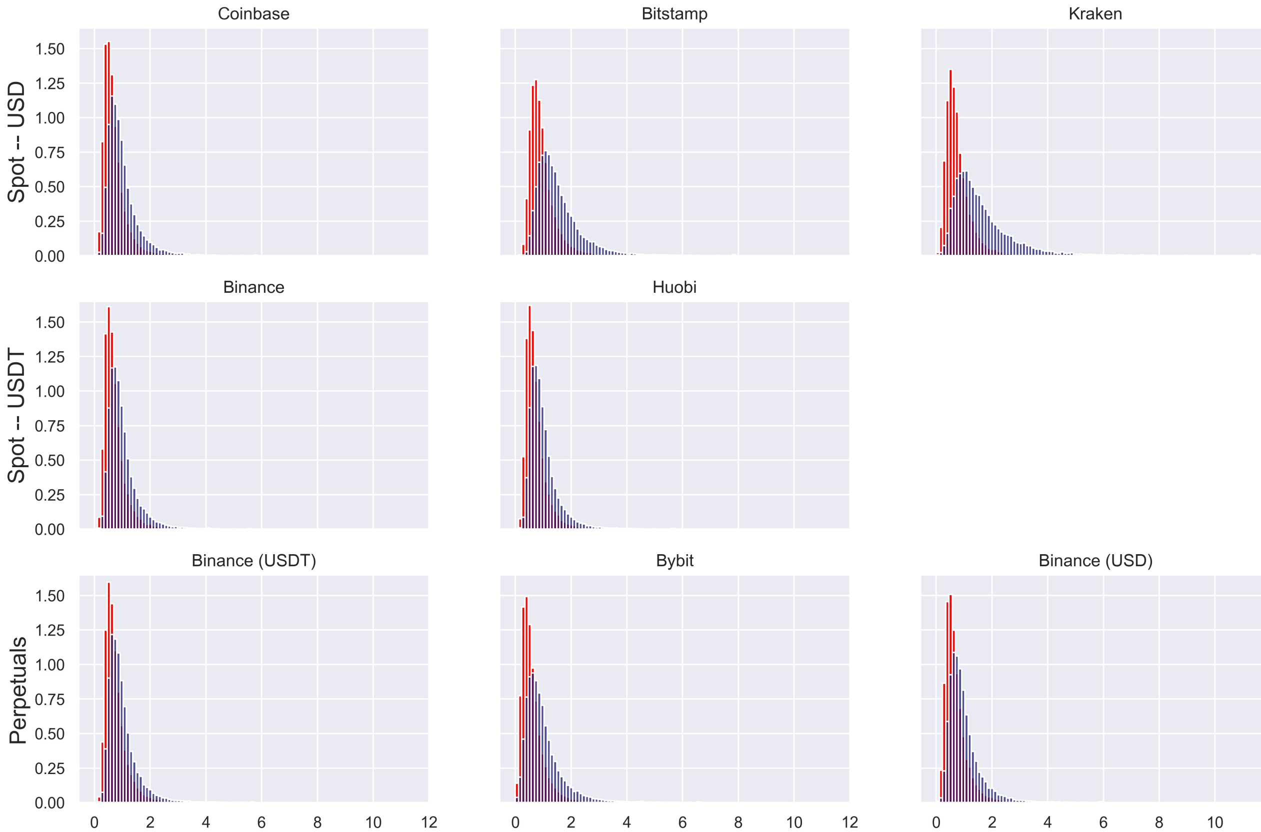

Note: The figure shows histograms of the five-minute realised volatility for USD spot pairs (top row), USDT spot pairs (middle row) and perpetuals (bottom row), both before (red) and after (blue) diurnal adjustment and winsorizing, covering the period from 1 January to 31 March 2021. The realised volatility is both adjusted for the deterministic intraday pattern using the average as diurnal adjustment factor and winsorized at the 99.95%-quantile. Note that due to the quite large number of zero observations on Kraken and Bitstamp, we excluded these zero values, allowing a better comparison across instruments.

Bayesian Information Criterion (BIC) and the half-life of the conditional mean. Overall, the volatility dynamics seem to be quite similar across all eight instruments. Not only is the total persistence of realised volatility – the sum of α and β – very high, varying between 0.96 and 0.98 for all instruments, but also the individual degrees of short-term persistence (captured by α) and long-term persistence (captured by β) are very similar at around $0.38-0.4$ and $0.57-0.6$, respectively. Only the Bybit contract slightly stands out – its short-term and long-term persistence are 0.29 and 0.68, respectively, which indicates that traders on Bybit are less reactive to prevailing market conditions than on the remaining seven instruments. A possible explanation for this might be the fee structure and the trading activity resulting from it. As stated above, Bybit offers maker rebates of 2.5bps for all users, allows trading with leverage 100 and its perpetual contract has a size of only $1. These properties make the product very attractive for smaller retail traders who are usually considered rather uninformed and therefore, the disclosure of new information or a change in market conditions has less influence on Bybit than on the other seven instruments.

Both the highest short-term persistence and the lowest long-term persistence are observed on the Binance USDT-perpetual. Interestingly, the two tether spot pairs (Binance, Huobi) exhibit the second and third highest (lowest) short-term (long-term) persistence. With standard errors of less than 0.01, the difference to the estimates for USD-instruments (Coinbase, Bitstamp, Kraken, Binance USD-perpetual) is not negligible. Therefore, we conclude that when updating their expectations, traders on the tether-based products – especially on the Binance USDT-perpetual – pay more attention to current market conditions than those on the USD-based instruments.

The half-life, which in our model measures the time it takes the logarithm of the conditional mean to halve its distance to the long-term mean, is quite similar for Coinbase, Huobi and the three Binance instruments. On these exchanges, it takes on average between 90 and 97 minutes until the effect of a volatility shock on the traders' expectation for the volatility within the next five-minute interval has halved. The lowest half-life is found on the Binance USDT-perpetual, which once more indicates that the traders on this product pay most attention to the current state of the market when updating their expectations. On Bitstamp, Kraken and Bybit, it takes the traders' expectation much longer to revert to the long-term average, which is probably caused by quite low trading volume on Bitstamp and Kraken and rather uninformed trading activity on Bybit.

As expected, the estimate of the auxiliary parameter α^0 is either not statistically significant or negative, therefore reducing the conditional mean after a zero observation occurs. Moreover, the estimate of the asymmetric response component is positive and highly significant on all eight instruments, which confirms presence of the leverage effect often documented in other asset classes. The strength of this effect is very similar across instruments at around 0.03, only Kraken stands out with a comparatively high estimate of 0.05. This implies that a negative return increases the volatility expected by traders for the next five-minute interval by about 7% to 10% more than a positive return. Finally, the distribution parameter s is quite similar on Coinbase, Bitstamp, Huobi and the three Binance instruments. Only Kraken and Bybit once again exhibit a significantly higher parameter estimate, indicating a greater standard deviation of innovations.

Compared to the findings of Nguyen et al. (2020) for the US Treasury market, all eight crypto instruments exhibit a much higher (lower) degree of short-term (long-term) persistence, implying higher sensitivity to prevailing market conditions and shorter memory than the Treasury notes. The total persistence however is higher on the US Treasury market, albeit only slightly (0.99 compared to between 0.96 and 0.98). Interestingly, during the 2007-2009 financial crisis as well as before and after economic announcements, the Treasury market exhibits elevated short-term persistence and its realised volatility dynamics appear to be more similar to those of crypto instruments.

5.2 Multivariate Dynamics

In this section, we analyse multivariate volatility flows among instruments. To get a detailed overview without overloading the number of parameters to be estimated, we fit the multivariate LogMEM_1 given in equation (3) for each group of instruments separately. That is, we first estimate the model for the three USD spot pairs (Coinbase, Bitstamp, Kraken), then we fit it to the two tether spot pairs (Binance, Huobi) and after that to the three perpetual swaps (Binance USDT, Bybit, Binance USD). Finally, we fit the model to all instruments jointly, except for Bitstamp and Kraken. Due to the comparatively large number of zero values, their exclusion makes estimation and interpretation of

Table 4 Univariate LogMEM$_1$(1,1)

	Coinbase	**Bitstamp**	**Kraken**	**Binance^S**	**Huobi**	**Binance^T**	**Bybit**	**Binance$^\$$**
ω	-0.0042^{**}	-0.1675^{***}	-0.1701^{***}	-0.0069^{***}	-0.0073^{***}	-0.0056^{***}	0.0025	-0.0042^{**}
α	0.3741^{***}	0.3837^{***}	0.3808^{***}	0.3896^{***}	0.3859^{***}	0.4012^{***}	0.2871^{***}	0.3781^{***}
α^0	0.0089	-0.0648^{***}	-0.2879^{***}	-0.1964^{*}	-0.0681	0.0020	0.0316	-0.2373^{***}
γ	0.0324^{***}	0.0379^{***}	0.0527^{***}	0.0326^{***}	0.0330^{***}	0.0305^{***}	0.0359^{***}	0.0345^{***}
β	0.5901^{***}	0.5995^{***}	0.5972^{***}	0.5733^{***}	0.5791^{***}	0.5610^{***}	0.6829^{***}	0.5851^{***}
s	0.2837^{***}	0.2826^{***}	0.4825^{***}	0.2617^{***}	0.2545^{***}	0.2548^{***}	0.4074^{***}	0.2958^{***}
LL	-753	$-21{,}217$	$-25{,}898$	503	1,314	1,448	$-8{,}353$	$-1{,}594$
BIC	1,568	42,494	51,857	-944	$-2{,}566$	$-2{,}835$	16,769	3,248
h	95	204	156	92	97	90	114	93

Note: The table reports parameter estimates, log-likelihood (LL), Bayesian Information Criterion (BIC) and half-life (h, in minutes) for the univariate LogMEM$_1$(1,1), fitted to diurnally-adjusted five-minute realised volatility over the period from 1 January to 31 March 2021. Binance^S, Binance^T and Binance$^\$$ represent the XBT/USD spot pair, the USDT-perpetual and the USD-perpetual, respectively, on Binance, while the parameter s denotes the standard deviations of the log residuals. The asterisks ***,**, * indicate significance at the 1%, 5%, and 10% level, respectively, based on robust standard errors. The half-life is calculated as $\log(0.5)/\log(\alpha+\beta)$, assuming strictly positive values and no asymmetric response.

the model more straightforward. To be consistent with the univariate analysis above, we choose to include only one lag for each model ($p = q = 1$).[16] In what follows, we document the results for the different instrument groups. Compared to the univariate analysis, the estimates for intercept, asymmetric response component and distribution parameter do not change significantly and are thus not reported. Similarly, we do not include the resulting estimates for matrix $\mathbf{A}^0$ since it contains purely auxiliary parameters and does not yield any further insights.

USD Spot Pairs: Table 5 reports the estimates for matrices $\mathbf{A}$ and $\mathbf{B}$ of the vLogMEM$_1$(1,1) fitted to diurnally-adjusted five-minute realised volatility on the three USD spot pairs. It also includes the model's log-likelihood values and BIC. We can see that the estimates on Coinbase change only marginally compared to the univariate results in table 4 – the short-term persistence α decreases from 0.37 to about 0.34, while the long-term persistence level β reduces only from 0.59 to 0.58 – indicating that the volatility on Coinbase is rather independent of Bitstamp and Kraken. In contrast, the levels of persistence on Bitstamp and Kraken change substantially once lagged volatility of Coinbase is included, which indicates a certain degree of dependence on the latter. On Bitstamp, short-term (long-term) persistence reduces from 0.38 (0.60) to 0.28 (0.55), while the estimates on Kraken drop from 0.38 and 0.60 to about 0.21 and 0.50.

The non-diagonal entries of matrix $\mathbf{A}$ which capture the lagged dependence among exchanges reveal strong volatility flows from Coinbase to Bitstamp and Kraken but not vice versa. While the estimates for spillovers from Coinbase to Bitstamp and Kraken are highly significant at 0.12 and 0.20, the reversed flows are much less (0.05 and 0.004, respectively) and in the case of Kraken, they are not even significant. Putting these numbers in simple economic terms, a 20%-shock in realised volatility on Coinbase (i.e. $\varepsilon_t = 1.2$) leads to an increase of about 2.2% (3.7%) in expected volatility

[16] If more lags are included, the overall results behave very similar to the univariate case and do not change significantly (see footnote 15).

on Bitstamp (Kraken).[17,18] On the other hand, a shock of the same magnitude on Bitstamp increases the expected volatility on Coinbase only by less than 1%. Among Bitstamp and Kraken, we detect similar unidirectional volatility flows. Spillovers from Bitstamp to Kraken are highly significant at 0.13, while the reversed flows are quite weak and only about 0.02. That is, a 20% volatility shock on Bitstamp leads to an increase of 2.4% in expected volatility on Kraken. In contrast, a 20%-shock on Kraken raises the expected volatility on Bitstamp only very marginally by 0.4%.

Following these results, we conclude that within the fiat-based spot market, Coinbase is the main source of volatility. It exhibits major flows to both Bitstamp and Kraken and receives only rather little volatility from Bitstamp. Among the two smaller exchanges, Bitstamp and Kraken, the former is the larger emitter of volatility, albeit its flows to Kraken are significantly smaller than those from Coinbase. Finally, Kraken transmits only very marginal flows to Bitstamp and is the main receiver of volatility.

Table 5 Multivariate $\text{LogMEM}_1(1,1)$ – USD Spot

		Coinbase	**Bitstamp**	**Kraken**
	Coinbase	0.3383	0.0546	0.0044^{ns}
A	Bitstamp	0.1184	0.2846	0.0248
	Kraken	0.2006	0.1285	0.2073
	Coinbase	0.5760		
B	Bitstamp		0.5522	
	Kraken			0.4975
LL		−653	−20,962	−25,418
BIC		1,408	42,025	50,937

Note: The table reports parameter estimates, log-likelihood value and BIC for the multivariate $\text{LogMEM}(1,1)_1$, fitted to five-minute realised volatility on Coinbase (CB), Bitstamp (BS) and Kraken (KK) over the period from 1 January to 31 March 2021. The superscript ^{ns} indicates that the estimate is not significant at the 1%-level. For better readibility, the parameters capturing short-term persistence are highlighted in red.

USDT Spot Pairs: Table 6 reports parameter estimates, log-likelihood values and BIC for the two tether spot pairs on Binance and Huobi. Similarly to the fiat-based analysis above, we see a significant reduction in short-term persistence on Binance – it drops from 0.39 in the univariate analysis to less than 0.26. Huobi on the other hand experiences only a slight reduction in short-term persistence from 0.39 to about 0.36, which suggests that realised volatility on Huobi is rather independendent of Binance. This is also confirmed by the non-diagonal entries of matrix **A**. With a highly significant estimate of 0.14, volatility flows from Huobi to Binance are very strong, whereas the reversed flows from Binance to Huobi are not even significant. Therefore, we conclude that within the tether-based spot market, volatility mainly emerges on Huobi from where it then spills over to Binance.

[17] Note that we use the log of the conditional mean in the LogMEM, therefore we first have to apply the exponential function to obtain the conditional mean itself.

[18] Also note that this amount is not the complete increase. Since the contemporaneous correlation among innovations, which our model does not capture, is still quite high even at the five-minute frequency, the actual increase will be a little more.

Table 6 Multivariate $\text{LogMEM}_1(1,1)$ – USDT Spot

		Binance^S	**Huobi**
A	Binance^S	0.2591	0.1377
	Huobi	0.0282^{ns}	0.3586
B	Binance^S	0.5679	
	Huobi		0.5780
LL		556	1,316
BIC		−1,030	−2,551

Note: The table reports parameter estimates, log-likelihood and BIC for the multivariate $\text{LogMEM}(1,1)_1$, fitted to five-minute realised volatility on Binance Spot (BI^S) and Huobi (HU) over the period from 1 January to 31 March 2021. The superscript ^{ns} indicates that the estimate is not significant at the 1%-level. For better readibility, the parameters capturing short-term persistence are highlighted in red.

Perpetual Swaps: Table 7 reports parameter estimates, log-likelihood value and BIC for the three perpetual swaps. We see that the dynamics of the Binance tether-margined perpetual are rather independent of the other two contracts, with its parameter estimates showing almost no changes compared to the univariate analysis. In contrast, the two fiat-based perpetuals on Binance and Bybit experience significant parameter changes – especially their short-term persistence reduces substantially from 0.29 to 0.14 for Bybit and from 0.38 to 0.25 for Binance. On Bybit, the long-term persistence also drops significantly to 0.61. These results suggest that realised volatility on the two USD-perpetuals on Binance and Bybit depends strongly on the tether-based Binance contract.

The non-diagonal entries of matrix **A** indicate strong unidirectional flows from the Binance USDT-perpetual to the two USD-based products. With estimates of 0.18 and 0.15, the Binance tether-contract transmits a great amount of volatility to both the Bybit and Binance fiat-perpetual – on Bybit, its influence even exceeds the short-term persistence. Put in economic terms once more, this means that a 20% volatility shock on the USDT-product leads to an increase of 3.3% (2.8%) in expected volatility on the Bybit (Binance) USD-contract. In contrast, the Binance tether-based perpetual does not receive any significant volatility flows, as can be seen from the two non-significant entries in the first row of **A**. Among the two USD-perpetuals, Binance seems to be more important in terms of volatility transmission. Spillovers from Binance to Bybit are estimated as 0.09, while the reversed flows are not even significant at the 1%-level.

Based on these results, we conclude that the main source of volatility within the perpetuals market is clearly the tether-margined contract on Binance. From there, large amounts of volatility are transmitted to the two largest fiat-based perpetuals on Bybit and Binance. In addition, there are smaller volatility flows from the Binance USD-contract to Bybit. The fact that Bybit only receives, but does not transmit volatility might again be related to its trader types. Both fees and contract specifications make the product highly attractive for retail traders, generating rather uninformed trading activity to which traders on Binance do not react. Moreover, it might increase the time that traders on Bybit need to pick up signals from other exchanges, leading to a lagged dependence on realised volatility of Binance.

Table 7 Multivariate LogMEM$_1$(1,1) – Perpetuals

		Binance^T	**Bybit**	**Binance$^\$$**
A	Binance^T	0.3939	0.0126^{ns}	-0.0089^{ns}
	Bybit	0.1831	0.1402	0.0867
	Binance$^\$$	0.1462	0.0083^{ns}	0.2471
B	Binance^T	0.5620		
	Bybit		0.6113	
	Binance$^\$$			0.5753
LL		1,450	−8,114	−1,549
BIC		−2,799	16,331	3,199

Note: The table reports parameter estimates, log-likelihood and BIC for the multivariate LogMEM(1,1)$_1$, fitted to five-minute realised volatility on the Binance USDT-perpetual (BI^T), Bybit perpetual (BY) and Binance USD-perpetual (BI$^\$$) over the period from 1 January to 31 March 2021. The superscript ^{ns} indicates that the estimate is not significant at the 1%-level. For better readibility, the parameters capturing short-term persistence are highlighted in red.

Main Instruments: Table 7 reports the estimates for matrices **A** and **B** of the six-dimensional LogMEM$_1$(1,1) fitted to five-minute realised volatility on Coinbase, Huobi, Bybit and the three Binance instruments. It also includes the model's log-likelihood values and BIC.[19] Once more, the degree of short-term persistence changes significantly compared to the univariate dynamics. The most extreme change occurs on the Binance USDT spot pair – the estimate reduces from 0.39 to 0.07 – while the Binance tether-perpetual is the only instrument where the inclusion of lagged interdependence leads to an increase in short-term persistence. As in the previous multivariate analyses, the remaining parameters (asymmetric response component, long-term persistence, standard deviation of log residuals) do not change significantly. Only on Bybit, the level of long-term persistence drops slightly from 0.68 in the univariate analysis to about 0.62.

Overall, the Binance USDT-perpetual has the second lowest level of long-term persistence – only on the Binance USDT spot pair, it is marginally smaller – and by far the highest degree of short-term persistence. Therefore, the result from the univariate analysis that traders on this contract are more sensitive to prevailing market conditions than on the remaining five instruments still holds, once we allow for lagged interdependence.

For better illustration, figure 6 shows the magnitude of volatility flows (i.e. the non-diagonal entries of matrix **A**) as a circular plot. Note that non-significant entries are set to zero. From this figure, it becomes clear that the Binance USDT-perpetual is the main source of volatility. It exhibits very strong spillovers to all other instruments, while it only receives quite weak volatility flows from Coinbase and Binance spot. The two remaining perpetuals (Bybit and Binance USD) are far less important and besides some volatility flows between the two products, there is only a very minor spillover from Bybit to Huobi. Interestingly, both Coinbase and Binance spot transmit volatility to all other instruments. However, as can be seen in table 8, all of these flows are negative. An explanation for this might be that traders on these spot exchanges need longer to react to (temporarily) increased volatility on the

[19] To detect potential time-variation, we have also estimated the model for each of the three months separately. The overall results however do not change significantly, indicating that the realised volatility dynamics are quite stable over time.

perpetual swaps. Once the volatility has risen on the spot exchanges, the transient increase on the perpetual contracts is already reversed, leading to a negative estimate of the volatility transmitted from spot exchanges. Finally, both Bybit and Huobi seem to be receivers rather than transmitters of volatility. Besides a weak spillover between the two instruments and minor flows from Huobi to Binance spot, we detect no significant influence on the remaining four instruments.

Table 8 Multivariate $\text{LogMEM}_1(1,1)$ – Main Instruments

		Coinbase	**Binance^S**	**Huobi**	**Binance^T**	**Bybit**	**Binance$^\$$**	To
A	Coinbase	0.2108	−0.0895	0.0352^{ns}	0.2548	0.0088^{ns}	-0.0173^{ns}	0.3762
	Binance^S	-0.0164^{ns}	0.0742	0.0626	0.2439	0.0105^{ns}	0.0223^{ns}	0.3806
	Huobi	−0.0282	−0.0851	0.2720	0.2177	0.0198	-0.0108^{ns}	0.3962
	Binance^T	−0.0307	−0.0995	0.0247^{ns}	0.4702	0.0149^{ns}	0.0132^{ns}	0.3400
	Bybit	−0.0761	−0.1373	0.0441^{ns}	0.3085	0.1450	0.1200	0.3600
	Binance$^\$$	−0.0594	−0.1039	0.0110^{ns}	0.2606	0.0134^{ns}	0.2742	0.3716
	From	−0.0163	−0.4411	0.3346	1.7557	0.1648	0.3943	
B	Coinbase	0.5749						
	Binance^S		0.5651					
	Huobi			0.5780				
	Binance^T				0.5665			
	Bybit					0.6167		
	Binance$^\$$						0.5805	
LL		−562	690	1,425	1,482	−8,072	−1,504	
BIC		1,287	−1,218	−2,687	−2,801	16,307	3,170	

Note: The table reports parameter estimates, log-likelihood and BIC for the multivariate $\text{LogMEM}(1,1)_1$, fitted to five-minute realised volatility on Coinbase (CB), Binance Spot (BI^S), Huobi (HU), Binance USDT-perpetual (BI^T), Bybit (BY) and Binance USD-perpetual ($\text{BI}^\$$) over the period from 1 January to 31 March 2021. The superscript ^{ns} indicates that the estimate is not significant at the 1%-level. For better readibility, the parameters capturing short-term persistence are highlighted in red.

5.3 Intraday Variation of Spillovers

So far, we have analysed volatility flows and co-movement over the complete course of the day. However, similar to the FX market, crypto exchanges operate globally and are open 24/7. As shown in figures 2 and 4, trading activity varies significantly over the course of the day and therefore, it is of particular importance to examine the volatility dynamics for short-term intraday variation.[20] We do so by decomposing the trading day into three different time periods of equal length, namely the Asian trading hours from midnight to 08:00 UTC, the European time period from 08:00 to 16:00 UTC and finally the US trading hours from 16:00 to midnight UTC. For each of these three time periods, we introduce in our multivariate $\text{LogMEM}_1(1,1)$ an interaction term of the log realised volatility with a dummy variable that is one during the respective time period and zero otherwise. Consequently, the

[20]Intraday time-variation of the long-term persistence is rather unlikely and therefore, we focus on variation in the short-term parameters.

Figure 6 Volatility Flows Among Main Instruments

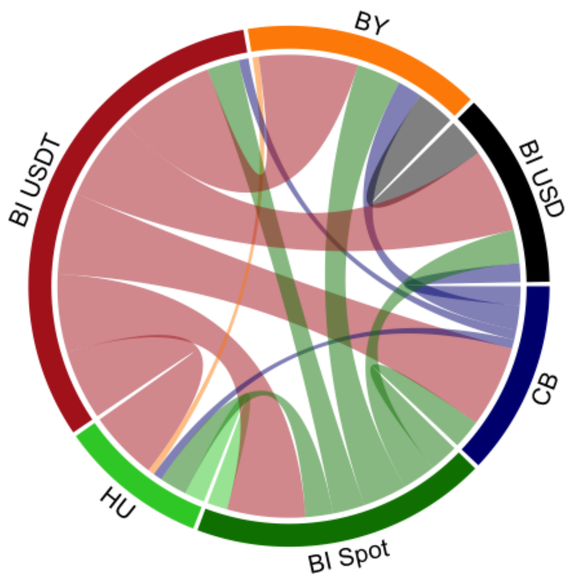

Note: The figure shows the magnitude of volatility flows among six major instruments, including Coinbase (CB), Binance spot (BI Spot), Huobi (HU), Binance USDT-perpetual (BI USDT), Bybit perpetual (BY) and Binance USD-perpetual (BI USD). The flows are estimated using a multivariate $\text{LogMEM}_1(1,1)$ on the five-minute realised volatility between 1 January and 31 March 2021. Both the diagonal of $\mathbf{A}$ and entries that are not significant are set to zero.

conditional mean specification in (3) changes to

$$\begin{aligned}\log \boldsymbol{\mu}_t = \mathbf{w} + (\mathbf{A} + \mathbf{A}_{\text{AS}}\mathbb{1}_{\text{AS}} + \mathbf{A}_{\text{EU}}\mathbb{1}_{\text{EU}} + \mathbf{A}_{\text{US}}\mathbb{1}_{\text{US}}) \left(\log \boldsymbol{x}_{t-1} \odot \mathbb{1}_{\{\boldsymbol{x}_{t-j}>0\}}\right) \\ + \mathbf{A}^0 \mathbb{1}_{\{\boldsymbol{x}_{t-1}=0\}} + \boldsymbol{\Gamma} \log \boldsymbol{x}_{t-1}^- + \mathbf{B} \log \boldsymbol{\mu}_{t-1}\end{aligned}$$

where $\mathbb{1}_{AS}$, $\mathbb{1}_{EU}$ and $\mathbb{1}_{US}$ denote the dummy variables for Asian, European and US trading hours, respectively. All other variables are the same as in (3).

Table 9 reports the resulting estimates for the short-term matrices $\mathbf{A}$, $\mathbf{A}_{\text{AS}}$, $\mathbf{A}_{\text{EU}}$ and $\mathbf{A}_{\text{US}}$, together with log-likelihood values and BIC. For reasons of clarity and comprehensibility, entries that are not significant at the 10%-level are already removed. We do not include the estimates for the remaining parameters since most of them do not change significantly – only the intercepts on Coinbase, Huobi and the three Binance instruments reduce slightly to about -0.05.

We can see that volatility flows from the spot exchanges (Coinbase, Binance^S, Huobi) are far more constant over the course of the day than those from the three perpetual swaps. In particular, flows transmitted by Coinbase do not exhibit any intraday variation and the estimates from matrix $\mathbf{A}$ are quite close to those in the previous analysis. Similarly, we do not find significant intraday variation on Huobi, where only two entries – the short-term persistence during European trading and the volatility flows to Binance spot during the US trading period – are negative and (weakly) significant. The Binance USDT spot pair on the other hand only transmits volatility during European trading hours and, to a lesser extent, during US trading times. The parameter estimates in the second column of $\mathbf{A}_{\mathrm{EU}}$ are very similar to the short-term estimates in table 8 and therefore, we conclude that the previoulsy documented volatility flows from the Binance USDT spot pair are mainly caused by European trading activity.

All three perpetual swaps exhibit a high degree of intraday variation, with only eight out of the 54 associated parameters of the interaction matrices $\mathbf{A}_{\mathrm{AS}}$, $\mathbf{A}_{\mathrm{EU}}$ and $\mathbf{A}_{\mathrm{US}}$ being not significant. As in the previous analysis, the Binance USDT-perpetual exhibits the strongest volatility flows to all other instruments. However, these flows intensify throughout the day and reach their maximum during US trading hours, when the spillovers to all other instruments are (highly) significant and (far) above 0.16. The Bybit and Binance USD-perpetuals both exhibit significant volatility spillovers over the whole day, but their intraday patterns differ from the tether-based contract. While the total volatility flows transmitted by Bybit are largest (lowest) in Asian (European) trading, the Binance USD-perpetual emits most volatility during European trading hours and only rather little in Asian trading. We conclude that even after accounting for intraday variation in volatility spillovers, the Binance tether-based contract is still the instrument that generates and transmits most of the volatility. However, its leadership role is less pronounced depending on the time of day and the Bybit perpetual in particular gains relative importance during Asian trading.

Overall, the volatility flows among the six crypto instruments seem to strengthen over the course of the day. While the total spillover amount – measured as the sum of absolute flows, excluding the short-term persistence parameter – is 1.78 in Asian trading, it increases by 30% to about 2.30 during European and US trading hours. Also, when considering only the short-term persistence (i.e. the diagonal entries of the matrices), we find a similar intraday pattern. In Asian and European trading, the total short-term persistence of the six instruments is about 1.43 and 1.50, respectively, while it amounts to more than 1.70 during US trading hours. That is, a volatility shock that occurs during the US trading period raises the traders' expectations on the realised volatility within the next five-minute interval significantly more than a shock during Asian or European trading times. These two findings suggest that (i) during US trading hours, traders pay more attention to prevailing market conditions when updating their expectations and (ii) the crypto market exhibits a higher interconnectedness when traditional Western stock markets are open, which is consistent with figure 2.

6 Conclusion

We analyse high-frequency realised volatility dynamics and spillovers in the bitcoin market during the most recent bull period from 1 January to 31 March 2021. Our analysis focuses on two (crypto)currency pairs, namely trading bitcoin against the US dollar (the main fiat-crypto pair) and trading bitcoin against tether (the main crypto-crypto pair). Based on second-by-second transaction data covering the major spot and perpetual exchanges, we estimate both univariate and multivariate versions of the Logarithmic Multiplicative Error Model – first for each group of instruments separately and then for the most important instruments jointly. Finally, we examine the inter-exchange spillovers for intraday variation by dividing the day into three different time zones of equal length (Asian, European, US

Table 9 Multivariate $\text{LogMEM}_1(1,1)$ – Intraday

		Coinbase	**Binance^S**	**Huobi**	**Binance^T**	**Bybit**	**Binance$^\$$**
A	Coinbase	0.1906***	–	0.0715*	–	–	0.1082***
	Binance^S	–	0.1226**	0.1119***	–	–	0.1046***
	Huobi	−0.0505**	–	0.3514***	–	–	0.0781**
	Binance^T	−0.0456**	–	0.0646*	0.2669***	–	0.1095***
	Bybit	−0.0934***	–	–	0.1671**	0.0435**	0.3117***
	Binance$^\$$	−0.0779***	–	–	0.1277**	–	0.3431***
$\mathbf{A}_{\text{AS}}$	Coinbase	–	–	–	0.1546**	0.0826***	−0.1902***
	Binance^S	–	–	–	0.1172**	0.0577***	−0.1087**
	Huobi	–	–	–	0.1325**	0.0658***	−0.1343***
	Binance^T	–	–	–	–	0.0720***	−0.1420***
	Bybit	–	–	–	–	0.2177***	−0.2702***
	Binance$^\$$	–	–	–	–	0.0729***	−0.1108*
$\mathbf{A}_{\text{EU}}$	Coinbase	–	−0.0978**	–	0.1813***	0.0546**	−0.1466***
	Binance^S	–	−0.0984*	–	0.1402**	0.0462**	–
	Huobi	–	–	−0.0689*	0.1701***	0.0486**	−0.1025*
	Binance^T	–	−0.1198***	–	0.1569***	0.0469**	−0.0855*
	Bybit	–	−0.1356***	–	0.1447**	0.1885***	−0.2818***
	Binance$^\$$	–	−0.1043**	–	–	–	–
$\mathbf{A}_{\text{US}}$	Coinbase	–	–	–	0.2221***	0.0781***	−0.1656***
	Binance^S	–	–	−0.0848**	0.1654***	0.0610***	−0.1073**
	Huobi	–	−0.0935*	–	0.2020***	0.0572***	−0.0956*
	Binance^T	–	–	–	0.1902***	0.0611***	−0.1148**
	Bybit	–	−0.1215**	–	0.1903***	0.1978***	−0.2793***
	Binance$^\$$	–	–	–	0.1257**	0.0627***	–
LL		−502	760	1,489	1,565	−8,012	−1,450
BIC		1,350	−1,175	−2,633	−2,785	16,370	3,246

Note: The table reports parameter estimates, log-likelihood and BIC for the multivariate $\text{LogMEM}(1,1)_1$, fitted to five-minute realised volatility on Coinbase (CB), Binance Spot (BI^S), Huobi (HU), Binance USDT-perpetual (BI^T), Bybit (BY) and Binance USD-perpetual ($\text{BI}^\$$) over the period from 1 January to 31 March 2021. The asterisks ***,**, * indicate significance at the 1%, 5%, and 10% level, respectively, based on robust standard errors. For better readibility, entries that are not significant at the 10%-level are already removed and the parameters capturing short-term persistence are highlighted in red.

trading hours).

First and foremost, we find that the tether-margined perpetual contract on Binance is clearly the main emitter of volatility. Throughout the day, it continuously transmits strong flows to all other instruments, including both USD and USDT spot pairs as well as USD-margined perpetual contracts. However, the strength of these volatility flows varies over the course of the day. They are lowest in Asian trading, then intensify substantially in European trading and finally peak during US trading hours. Out of all instruments, the tether-based perpetual contract also receives the lowest volatility flows, which apart from one minor exception holds for all three time zones. Moreover, the tether-based contract generally has the highest short-term persistence and the second lowest long-term persistence, which indicates that traders on this contract are more reactive to prevailing market conditions than on the remaining instruments. However, as for the volatility flows, the short-term persistence varies over the course of the day and is significantly higher during US trading hours than in Asian trading.

From the remaining crypto instruments included in our analysis, much weaker volatility flows emanate and these products are thus far less important than the tether-contract in terms of volatility transmission. While we find no to very little intraday variation in spillovers from Coinbase and Huobi spot trading, the Binance USDT spot pair mainly transmits volatility during European trading hours. The two USD-margined perpetuals from Bybit and Binance exhibit strong intraday variation, but in contrast to the tether-based contract, volatility flows from these two products are strongest during Asian and European trading hours, respectively.

These results on their own are not particularly surprising or alarming. However, if they are combined with tether's continuously increasing market capitalization – at the time of writing, it was the third largest cryptocurrency, with a circulating supply of almost 63 bn tokens – its growing usage as both quote currency for many other cryptocurrencies and margin and settlemtn currency for the largest (and possibly highly-leveraged) cryptocurrency derivatives, the recurring allegations of misconduct and price inflation against its issuer, paired with their unwillingness to undergo a proper audit, it becomes clear that tether deserves and requires more attention. Academics and regulators in particular should expand their focus from bitcoin-fiat pairs to bitcoin-tether products. As we have shown, these are the main source of volatility within the bitcoin market and therefore have the highest potential for contagion.

Overall, the volatility flows among the six crypto instruments strengthen over the course of the day. In European and US trading, the total spillover amount is about 30% higher than during Asian trading hours. Moreover, the combined short-term persistence of all instruments is more than 14% higher during the US trading period than in Asian and European trading. This leads to the conclusion that (i) during US trading hours, traders pay more attention and are more reactive to prevailing market conditions when updating their expectations and (ii) the crypto market exhibits a higher interconnectedness when traditional Western stock markets are open.

For all eight instruments considered here, we identify presence of the leverage effect which is often documented in other asset classes (Bollerslev et al., 2006). The effect is quite strong, with a negative shock increasing the expected volatility by 7% – 10% more than a positive shock. Compared to the US Treasury market (Nguyen et al., 2020), the realised volatility of crypto instruments exhibits a similarly high degree of total persistence. However, the short-term (long-term) persistence is much higher (lower) than for the US Treasury notes, implying a higher sensitivity of the crypto products to prevailing market conditions. Only during the 2007-2009 financial crisis as well as before and after economic announcements, the realised volatility dynamics on the Treasury market seem to be similar to those of crypto instruments.

As a by-product to our spillover analysis, we document a very similar intraday pattern in trading volume and realised volatility for the eight crypto instruments. Between midnight and 16:00 UTC, they evolve in a U-shape, reaching their minimum during early UTC morning. Afterwards, trading volume and realised volatility on all instruments slowly decay to an average level. In fact, this pattern seems quite similar to FX markets, as documented by Andersen and Bollerslev (1997). Finally, we also observe distinct volume and volatility spikes every four hours, but they are most extreme at midnight and 16:00 UTC. We think that these spikes are related to the funding of perpetual contracts which happens every four hours (on different products). Possibly, some market participants modify their positions across multiple exchanges just after funding on some perpetual has occurred or they take advantage of some mispricing, either between different perpetuals or between spot and perpetuals.

In terms of future research, it might be interesting to assess information processing and volatility transmission from the different perspective of volatility discovery. Using the methodology proposed in Dias et al. (2016) and applied to three spot exchanges by Dimpfl and Elshiaty (2021), one could examine how much the different crypto instruments contribute to the common latent volatility process and whether the tether-margined perpetual contract on Binance – which we identify as the main source of volatility – is also the leading market in the volatility process. Moreover, due to the recently declining importance of bitcoin within the crypto market, a similar volatility transmission analysis within the largest altcoins, in particular ether, or even across the most important cryptocurrencies – including both spot and perpetuals – may be an interesting research topic.

References

Allen, D., F. Chan, M. McAleer, and S. Peiris (2014). Finite sample properties of the QMLE for the Log-ACD model: Application to Australian stocks. *Journal of Financial Econometrics 12*(1), 89–121.

Andersen, T. G. (2000). Some reflections on analysis of high-frequency data. *Journal of Business & Economic Statistics 18*(2), 146–153.

Andersen, T. G. and T. Bollerslev (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance 4*(2-3), 115–118.

Baur, D. G. and L. T. Hoang (2021). A crypto safe haven against bitcoin. *Finance Research Letters 38*.

Bauwens, L. and P. Giot (2000). The logarithmic ACD model: An application to the bid-ask quote process of three NYSE stocks. *Annales d'Économie et de Statistique 60*, 117–149.

Bollerslev, T., J. Litvinova, and G. Tauchen (2006). Leverage and volatility feedback effects in high-frequency data. *Journal of Financial Econometrics 4*(3), 353–384.

Cai, C. X., R. Hudson, and K. Keasey (2004). Intra day bid-ask spreads, trading volume and volatility: Recent empirical evidence from the London Stock Exchange. *Journal of Business Finance and Accounting 31*(5-6), 647–676.

Caporale, G. M., W.-Y. Kang, F. Spagnolo, and N. Spagnolo (2021). Cyber-attacks, spillovers and contagion in the cryptocurrency markets. *Journal of International Financial Markets, Institutions & Money*.

Cheah, E.-T., T. Mishra, M. Parhi, and Z. Zhang (2018). Long memory interdependency and inefficiency in bitcoin markets. *Economics Letters 167*, 18–25.

Cipollini, F., R. F. Engle, and G. M. Gallo (2017). Copula–based vMEM specifications versus alternatives: The case of trading activity. *Econometrics 5*(2).

Danielsson, J. and R. Payne (2001). Measuring and explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2. *Working Paper*.

Dias, G. F., C. M. Scherrer, and F. Papailias (2016). Volatility discovery. *Working Paper*.

Dimpfl, T. and D. Elshiaty (2021). Volatility discovery in cryptocurrency markets. *Working Paper*.

Engle, R. F. (2002). New frontiers for arch models. *Journal of Applied Econometrics 17*(5), 425–446.

Engle, R. F. and J. R. Russell (1998). Autoregressive Conditional Duration: A new model for irregularly spaced transaction data. *Econometrica 66*(5), 1127–1162.

Eross, A., F. McGroarty, A. Urquhart, and S. Wolfe (2019). The intraday dynamics of bitcoin. *Research in International Business and Finance 49*, 71–81.

Escribano, A. and G. Sucarrat (2018). Equation-by-equation estimation of multivariate periodic electricity price volatility. *Energy Economics 74*, 287–298.

Griffin, J. and A. Shams (2020). Is bitcoin really untethered? *Journal of Finance 75*(4), 1913–1964.

Hautsch, N., P. Malec, and M. Schienle (2014). Capturing the zero: A new class of zero-augmented distributions and multiplicative error processes. *Journal of Financial Econometrics 12*(1), 89–121.

Hautsch, N. and M. Podolskij (2013). Preaveraging-based estimation of quadratic variation in the presence of noise and jumps: Theory, implementation, and empirical evidence. *Journal of Business & Economic Statistics 31*(2), 165–183.

Jacod, J., Y. Li, P. A. Mykland, M. Podolskij, and M. Vetter (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications 119*, 2249–2276.

Jain, P. C. and G.-H. Joh (1988). The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis 23*(3), 269–283.

Ji, Q., E. Bouri, L. Kristoufek, and B. Lucey (2021). Realised volatility connectedness among bitcoin exchange markets. *Finance Research Letters 38*.

Katsiampa, P., S. Corbet, and B. Lucey (2019). High frequency volatility co-movements in cryptocurrency markets. *Journal of International Financial Markets, Institutions & Money 62*, 35–52.

McGroarty, F., O. ap Gwilym, and S. Thomas (2009). The role of private information in return volatility, bid–askspreads and price levels in the foreign exchange market. *Journal of International Financial Markets, Institutions and Money 19*, 387–401.

McInish, T. H. and R. A. Wood (1990). An analysis of transactions data for the Toronto Stock Exchange: Return patterns and end-of-the-day effect. *Journal of Banking and Finance 14*(2-3), 441–458.

Nguyen, G., R. Engle, M. Fleming, and E. Ghysels (2020). Liquidity and volatility in the U.S. treasury market. *Journal of Econometrics 217*, 207–229.

Ozenbas, D. (2008). Intra-day trading volume patterns of equity markets: A study of US and European stock markets. *International Business and Economics Research Journal 7*(8).

Sensoy, A., T. C. Silva, S. Corbet, and B. M. Tabak (2021). High-frequency return and volatility spillovers among cryptocurrencies. *Applied Economics*.

Taylor, N. and Y. Xu (2017). The logarithmic vector multiplicative error model: an application to high frequency NYSE stock data. *Quantitative Finance 17*(7), 1021–1035.

Wang, J. and G. M. Ngene (2020). Does bitcoin still own the dominant power? an intraday analysis. *International Review of Financial Analysis 71*.

Yi, S., Z. Xu, and G.-J. Wang (2018). Volatility connectedness in the cryptocurrency market: Is bitcoin a dominant cryptocurrency? *International Review of Financial Analysis 60*, 98–114.