Volatility Cascades in Cryptocurrency Trading*

Nikola Gradojevic
University of Guelph
ngradoje@uoguelph.ca

Ilias Tsiakas
University of Guelph
itsiakas@uoguelph.ca

March 2021

Abstract

This paper studies volatility cascades across multiple trading horizons in cryptocurrency markets. Using one-minute data on Bitcoin, Ethereum and Ripple against the US dollar, we implement the wavelet Hidden Markov Tree model. This model allows us to estimate the transition probability of high or low volatility at one time scale (horizon) propagating to high or low volatility at the next time scale. We find that when moving from long to short horizons, volatility cascades tend to be symmetric: low volatility at long horizons is likely to be followed by low volatility at short horizons, and high volatility is likely to be followed by high volatility. In contrast, when moving from short to long horizons, volatility cascades are strongly asymmetric: high volatility at short horizons is now likely to be followed by low volatility at long horizons. These results are robust across time periods and cryptocurrencies.

Keywords: Cryptocurrencies; Bitcoin; Ethereum; Ripple; Volatility Cascade; Wavelet Hidden Markov Tree model.

JEL classification: F31; G15; C58.

**Corresponding author*: Nikola Gradojevic, Department of Economics and Finance, Lang School of Business and Economics, Univer-sity of Guelph, Guelph, Ontario N1G 2W1, Canada. Tel: (+1) 519-824-4120. Email: ngradoje@uoguelph.ca.

1. Introduction

Market participants with a different time horizon observe a different component of the volatility of financial asset returns (Mandelbrot, 1963, 1997). For example, daily volatility and monthly volatility are distinct volatility components with distinct properties (Corsi, 2009). Volatility components exhibit clustering across time scales. In other words, there are clusters of high and low volatility associated with both short and long horizons (Calvet and Fisher, 2002; Bacry *et al.*, 2008).[1] In this context, a volatility cascade is the transmission of volatility clustering across time scales. As shown by Gençay *et al.* (2010), volatility cascades tend to be asymmetric: low volatility at a low-frequency scale is typically followed by low volatility at all higher frequency scales. However, the reverse does not hold: high volatility at a low-frequency scale is not typically followed by high volatility at higher frequency scales. As a result, long-term volatility is a good predictor of short-term volatility when volatility is low, but not when volatility is high.

In this paper, we study volatility cascades in cryptocurrency markets. We do so by addressing two related questions. First, is there a transmission of volatility from long to short horizons (or vice versa)? Second, does the transmission of volatility depend on the state of volatility (i.e., low vs. high volatility)? For example, if long-horizon volatility is high will short-horizon volatility be high as well? In other words, does low or high volatility at one time scale transmit to low or high volatility at the next time scale? These questions allow us to determine whether there is a volatility cascade from long horizons (low-frequency) to short horizons (high-frequency) or vice versa. They also provide a framework for assessing whether the transmission of volatility is asymmetric across time horizons. In short, the main objective of this paper is to provide an examination of the transmission of different states of cryptocurrency volatility across investor horizons.

[1] Volatility clustering across time scales is typically described by multiscaling: volatility varies across time scales as a power law of the time increment. Multiscaling is also referred to as multifractality. Volatility clustering at all frequencies is consistent with the intuition that economic factors such as technological shocks, business cycles and liquidity shocks have different time scales (Calvet and Fisher, 2002).

Our empirical analysis focuses on three cryptocurrency exchange rates: Bitcoin, Ethereum and Ripple relative to the US dollar. These represent the most popular cryptocurrencies in worldwide trading.[2] We employ a high-frequency data set comprised of one-minute returns for the sample period of January 1, 2015 to October 15, 2020. Armed with these data, we classify returns and volatility into different time horizons (scales). The time scales range from one minute to about one year. This provides us with a wealth of information on the full range of horizons for high-frequency and low-frequency traders.

Our empirical methodology is based on the wavelet Hidden Markov Tree (HMT) model as implemented by Gençay *et al.* (2010). In this model, each time scale is characterized by a two-state regime of high and low volatility. Overall, the wavelet HMT model is designed to capture the dependency between high-volatility states and low-volatility states at each point in time across multiple time scales. Specifically, the model separates each time-scale component from the observed data using the wavelet coefficients. Then, it estimates the transition probability of moving from a high-volatility or low-volatility state at one time scale to a high-volatility or low-volatility state at the next time scale. The transition probabilities enable an assessment of whether there exists a symmetric or asymmetric vertical dependence in cryptocurrency volatility. For example, is high volatility at long horizons followed by high or low volatility at short horizons? It is important to note that our analysis implements a bi-directional wavelet HMT model, where information from longer time horizons directly influences shorter time horizons and vice versa. Results for the two cases are reported separately.

Our main empirical finding is that for cryptocurrencies, when we move from long to short horizons, volatility cascades tend to be mostly symmetric. However, when we move from short to long horizons, volatility cascades are strongly asymmetric. In particular, consider the first case of moving from long to short horizons (i.e., from a low-frequency time

[2] As of December 2020, Bitcoin and Ethereum are the top two cryptocurrencies by market value, while Ripple is the fourth. Tether is the third by market value, but its value is pegged to the US dollar so that the Tether/US dollar exchange rate is always equal to one. For this reason, Tether is not part of our analysis.

scale to a high-frequency time scale). Then, according to the transition probabilities of the wavelet HMT model, low volatility is very likely to be followed by low volatility across time scales; and high volatility is also likely to be followed by high volatility across time scales. This persistence in volatility transmission defines symmetric vertical dependence across time scales.

Next, consider the second case of moving from short to long horizons (i.e., from a high-frequency time scale to a low-frequency time scale). Low volatility is still very likely to be followed by low volatility across time scales. However, high volatility is now likely to be followed by low volatility across time scales. In other words, high volatility at short horizons does not lead to high volatility at longer horizons. This lack of persistence in the transmission of high volatility defines asymmetric vertical dependence across time scales.

In this context, our paper makes a threefold contribution to the literature. First, to our best knowledge, this is the first paper to study volatility cascades in the emerging cryptocurrency markets using high-frequency data. Second, we implement a bi-directional approach to volatility cascades using the wavelet HMT model, which is a novel contribution to the literature since earlier work on other markets typically focuses on one-direction volatility cascades. Third, we show that earlier empirical results indicating the presence of asymmetric vertical dependence in the volatility transmission of standard exchange rates do not always apply to cryptocurrency markets. Our findings indicate that vertical dependence across time scales is symmetric when moving from long to short horizons but strongly asymmetric when moving from short to long horizons.

The emphasis on different time scales is motivated by the fact that different types of traders have different horizons. For example, long-term traders trade infrequently and typically react to information on economic fundamentals that arrives at low frequency. In contrast, short-term traders trade frequently, often many times a day, and typically follow asset prices at high-frequency. The combination of trading activities across trader types classified by their horizon generates the observed market prices.

Trader heterogeneity across time horizons is formally studied by Müller *et al.* (1997), who propose the Heterogeneous Market Hypothesis. The main idea of the hypothesis is that traders with different time horizons perceive, react to and cause different types of volatility. The result of trader heterogeneity is a volatility cascade from the low frequencies to the high frequencies. The economic mechanism is as follows: short-term traders react to changes in long-term volatility because it may affect their long-term expected risk and return. This reaction causes short-term volatility. In contrast, short-term volatility does not affect the trading activity of long-term traders. Hence, short-term volatility does not cause long-term volatility. This hierarchical structure indicates that high long-term volatility will cause high short-term volatility (i.e., symmetric vertical dependence) but high short-term volatility will not cause high long-term volatility (i.e., asymmetric vertical dependence). In short, therefore, our findings are consistent with the Heterogeneous Market Hypothesis of Müller *et al.* (1997).[3]

The meteoric rise of cryptocurrencies as a new investment asset class has triggered an emerging literature on the statistical behavior of cryptocurrency returns and their investment diversification gains. For example, one line of research assesses predictability in cryptocurrency returns (see, e.g., Panagiotidis *et al.*, 2018, Adcock and Gradojevic, 2019, and Bouri *et al.*, 2019). A second line of research demonstrates that the volatility of cryptocurrency returns exhibits clustering and long memory (see, e.g., Dyhrberg, 2016, Katsiampa, 2017, Klein *et al.*, 2018, Ardia *et al.*, 2019, Hafner, 2020, and Segnon and Bekiros, 2020). A third line of research focuses on jumps in both returns and volatility (see, e.g., Gronwald, 2019, Bouri *et al.*, 2020, and Scaillet *et al.*, 2020). A fourth line of research indicates that cryptocurrencies exhibit high correlation with one another but low correlation with traditional assets such as stocks, bonds, commodities and hard currencies (see, e.g., Borri, 2019). These properties make cryptocurrencies valuable to investors for their portfolio diversifica-

[3]To illustrate this idea, consider a simple example: a quick price increase in an asset followed by an immediate decrease might be an important event for a high-frequency trader, but it is a non-event for a low-frequency trader. Conversely, if the change in price lasts a few months because of a shock in economic fundamentals, then this is an important event for both short-term and long-term traders.

tion benefits (see, e.g., Briere *et al.*, 2015, Eisl *et al.*, 2020, Hu *et al.*, 2019, Charfeddine *et al.*, 2020, and Trimborn *et al.*, 2020). A fifth line of research studies the cross-section of cryptocurrency returns to identify common risk factors based on cryptocurrency market return, market capitalization and momentum (Bhambhwani *et al.*, 2020; Borri and Shakhnov, 2020a; Liu *et al.*, 2020). A sixth line of research demonstrates that regulatory actions such as trading restrictions can have a strong impact on cryptocurrency markets, both nationally and internationally (Auer and Claessens, 2018; Borri and Shakhnov, 2020b). Finally, other relevant studies include Makarov and Schoar (2020) for a study of price formation across cryptocurrency exchanges; Easley *et al.* (2019) for a study of Bitcoin transaction fees; and Trimborn and Hardle (2018) for proposing a new cryptocurrency index called CRIX.

This paper is related to Celeste *et al.* (2020), which is a study of multifractality in the volatility of cryptocurrency returns. Using daily returns on the same cryptocurrencies, Bitcoin, Ethereum and Ripple, Celeste *et al.* (2020) implement Continuous Wavelet Transform analysis to study price dynamics in the time as well as the frequency domain. Their main finding is that the price dynamics of cryptocurrencies are dominated by short-term speculators, as opposed to buy-and-hold investors. The analysis of Celeste *et al.* (2020), however, substantially deviates from our paper in a number of ways: (1) it does not consider high-frequency investment horizons as we do; and (2) the empirical approach is not based on the wavelet-based HMT of Gençay *et al.* (2010) that we implement. As a result, Celeste *et al.* (2020) do not explicitly identify high-volatility and low-volatility regimes, and the associated volatility cascades from high-frequency to low-frequency scales and vice versa. In short, therefore, our paper makes a distinct contribution to the rapidly expanding literature on the volatility of cryptocurrency returns. For other applications of wavelet analysis to Bitcoin returns and volatility, see Kristoufek (2015) and Kang *et al.* (2019).

The paper is organized as follows. In the next section, we describe the methodology for implementing the wavelet HMT model and estimating the transition probabilities. The high-frequency data on the three cryptocurrencies are analyzed in Section 3. Section 4 reports and discusses the empirical results. Finally, Section 5 concludes.

2. Methodology

Our methodology implements the approach of Gençay *et al.* (2010), which is based on a combination of the Discrete Wavelet Transform (DWT) together with the Hidden Markov Tree (HMT) model in what is known as the wavelet HMT model. The first component of the wavelet HMT model is the DWT, which is designed to capture information at different frequencies at each point in time. Specifically, the DWT produces a set of wavelet coefficients, which decompose the information from the original time series into pieces associated with both time and frequency.

It is important to note that, by design, the DWT is applied to a dyadic vector of observations N, defined as $N = 2^j$, where $j = 1, \ldots, J$ is a positive integer denoting the time scale. Our analysis is based on $J = 19$ time scales. In this framework, $J = 0$ corresponds to the original one-minute returns, $J = 1$ corresponds to the first time scale ranging from 2-4 minutes, $J = 2$ ranges from 4-8 minutes and so on. The last time scale is $J = 19$, which corresponds to a range of about 1-2 years. Table 1 presents the time horizons that each of the 19 time scales correspond to. Overall, wavelet multiscaling preserves the information from the original one-minute returns, while allowing us to investigate returns and volatility at multiple time horizons simultaneously.

The second component of the model is HMT, which allows small or large wavelet coefficients to be connected vertically across scales. In other words, the HMT model allows for wavelet dependence across scales but not within scales. In our analysis, we implement a wavelet HMT model, which classifies volatility at each scale in two states $s \in \{0, 1\}$: the low-volatility state denoted by $s = 0$ and the high-volatility state denoted by $s = 1$. Then, the transition probability of volatility from one scale to the next is defined as follows.

Definition 1 *The transition probability* $p_{s,j}$ *is the likelihood of observing a low (high) volatil-*

Scale (j)	2^j	Time Horizons				
		Minutes	Hours	Days	Months	Years
1	2	$2-4$				
2	4	$4-8$				
3	8	$8-16$				
4	16	$16-32$				
5	32	$32-64$	$0.5-1.1$			
6	64		$1.1-2.1$			
7	128		$2.1-4.3$			
8	256		$4.3-8.5$			
9	512		$8.5-17.1$			
10	1,024		$17.1-34.1$	$0.7-1.4$		
11	2,048			$1.4-2.8$		
12	4,096			$2.8-5.7$		
13	8,192			$5.7-11.4$		
14	16,384			$11.4-22.8$		
15	32,768			$22.8-45.5$	$0.8-1.5$	
16	65,536				$1.5-3.0$	
17	131,072				$3.0-6.1$	
18	262,144				$6.1-12.1$	$0.5-1.0$
19	524,288					$1.0-2.0$

Table 1: Time Scales

The table presents a translation of wavelet scales into time horizons. Each scale corresponds to a frequency interval and, therefore, each time scale is associated with a range of time horizons.

ity state s at time scale j provided that there is a low (high) volatility state s at time scale j+1.

In other words, the transition probability reflects the likelihood of transmitting low or high volatility from one time scale to the next. Note that the transition probability will be assessed bi-directionally: (1) when moving from a long to a short horizon, i.e., moving to scale j given the volatility state at scale $j+1$ (lower frequency); and (2) moving from a short to a long horizon, i.e., moving to scale $j+1$ given the volatility state at scale j (higher frequency). The results for the two cases will be reported separately. The notation in this section is based on the first case of moving from a long to a short horizon.[4]

Based on Definition 1, the wavelet HMT model assumes that the transition probability

[4]Note that our bi-directional model is an extension of the original model of Gençay *et al.* (2010), which was one-directional only allowing information from longer horizons to influence shorter horizons.

matrix of volatility is scale-dependent and has the following form:

$$A_j = \begin{bmatrix} p_{0,j} & 1 - p_{0,j} \\ 1 - p_{1,j} & p_{1,j} \end{bmatrix}, \quad \text{for } j = 1, \ldots, J - 1,$$

where the *conditional* transition probabilities are defined as follows:

$$\left\{ \begin{aligned} p_{0,j} &= P(\text{low volatility at scale } j \mid \text{low volatility at scale } j+1) \\ 1 - p_{0,j} &= P(\text{high volatility at scale } j \mid \text{low volatility at scale } j+1) \\ p_{1,j} &= P(\text{high volatility at scale } j \mid \text{high volatility at scale } j+1) \\ 1 - p_{1,j} &= P(\text{low volatility at scale } j \mid \text{high volatility at scale } j+1) \end{aligned} \right\}. \tag{1}$$

The conditional probabilities reflect the persistence of volatility across time scales. Consider, for example, probability $p_{1,j}$. This probability addresses the following question: given a high-volatility state at scale $j+1$, how likely is the high-volatility state to persist to scale j? In other words, when moving from long to short horizons, if long-horizon investors face high volatility, does this also trigger high volatility for short-horizon investors? Conversely, when moving from short to long horizons, if short-horizon investors face high volatility, does this also trigger high volatility for long-horizon investors?

In our empirical analysis, we refer to $p_{0,j}$ as the "low-to-low" probability, $1 - p_{0,j}$ as the "low-to-high" probability, $p_{1,j}$ as the "high-to-high" probability and, finally, $1 - p_{1,j}$ as the "high-to-low" probability.

Following Gençay *et al.* (2010), the wavelet HMT is modelled as a two-state mixture distribution for the two volatility states $s \in \{0, 1\}$. The two components of the mixture distribution are both Gaussian probability density functions (PDFs), where the first Gaussian PDF has moments (μ_0, σ_0^2) and the second Gaussian PDF has moments (μ_1, σ_1^2) such that $\sigma_1^2 > \sigma_0^2$. It is well known that a mixture of two or more Gaussian distributions with different variances is typically not Gaussian.

Given the properties of the wavelet HMT model, as described in Gençay *et al.* (2010), the full likelihood is defined as follows:

$$
\begin{aligned}
f_{\mathbf{W}}(\mathbf{W}) \;=\; \sum_{\mathbf{S}} \Bigg\{ & P(S_{V_J} = s_{V_J}) f_{V_J|S_{V_J}}(v_J)\, P(S_{J,0} = s_{J,0}) f_{W_{J,0}|S_{J,0}}(w_{J,0}) \\
& \times \prod_{j=1}^{J-1} \prod_{n=0}^{N/2^j-1} f_{W_{j,n}|S_{j,n}}(w_{j,n}) P(S_{j,n} = s_{j,n}|S_{j+1,\lfloor n/2 \rfloor}) \Bigg\},
\end{aligned} \tag{2}
$$

where $\mathbf{W}$ is the vector of wavelet coefficients with elements $W_{j,n}$, j is the scale, $n = 0, 1, \ldots, 2^j - 1$, and $S_{j,n}$ is the (hidden) random volatility state variable with two possible values $s \in \{0, 1\}$.

For this wavelet HMT specification, the complete parameter vector is as follows:

$$
\theta = (p_{0,1}, p_{1,1}, \ldots, p_{0,J}, p_{1,J}, \mu_{0,1}, \sigma^2_{0,1}, \ldots, \mu_{0,J}, \sigma^2_{0,J} \mu_{1,1}, \sigma^2_{1,1}, \ldots, \mu_{1,J}, \sigma^2_{1,J}), \tag{3}
$$

where $p_{i,j} = p_{0,1}, p_{1,1}, \ldots, p_{0,J}, p_{1,J}$ denotes the transition probabilities and (μ, σ^2) are the parameters of the Gaussian PDFs. The DWT ensures that the expected value of all wavelet coefficients is zero when using a wavelet filter of sufficient length. We thus make the assumption that $\mu_{s,j} = 0$ for all s and j.

The final step to our methodology involves estimation of the wavelet HMT model, which is performed using the Expected Maximization (EM) algorithm. The EM algorithm jointly estimates the parameters θ and the distribution of the hidden volatility states given the observed wavelet coefficients, $\mathbf{W}$. Specifically, we calculate the log-likelihood of the wavelet HMT in Equation (2) using an upward-downward algorithm (see, e.g., Crouse *et al.*, 1998). This produces estimates of the model parameters, and particularly the transition probabilities, which are the focus of our empirical analysis.

3. Cryptocurrency Data

Our empirical analysis is based on high-frequency data for three popular cryptocurrencies. Specifically, we use one-minute returns for three cryptocurrency exchange rates relative to the US dollar: Bitcoin (BTC/USD), Ethereum (ETH/USD), and Ripple (XRP/USD). The data span the following sample periods: January 1, 2015 to October 15, 2020 for BTC/USD; January 1, 2016 to October 15, 2020 for ETH/USD; and January 1, 2018 to October 15, 2020 for XRP/USD. The different starting dates of the sample periods are due to data availability. Our data include all weekends and holidays since we find that these days are active in cryptocurrency markets and high-frequency data are hence available. All data are obtained from *CryptoCompare.com*.[5]

We define the one-minute return on each exchange rate as $r_t = ln(P_t) - ln(P_{t-1})$, where P_t is the one-minute exchange rate provided by CryptoCompare.com. The volatility of the cryptocurrency exchange rate returns over a given horizon is defined as the sum of squared one-minute returns r_t^2 over the horizon (see, e.g., Andersen *et al.*, 2001).

To illustrate the data used in our analysis, we report the daily price and volatility for BTC/USD in Figure 1, for ETH/USD in Figure 2, and for XRP/USD in Figure 3. Figure 1 indicates that BTC/USD experienced exponential growth in 2017 and a drastic decline in 2018, whereas the COVID-19 global pandemic in 2020 is associated with the highest volatility. The same pattern is observed for ETH/USD. XRP/USD exhibits different behaviour as its price has been declining since the beginning of our data sample in 2018. Moreover, XRP/USD displays high volatility in both 2018 and 2020.

As explained in the previous section, wavelet analysis requires a dyadic sample with sample size $N = 2^j$, for $j = 1, \ldots, J$. We select $J = 19$ to implement the 19 time scales illustrated in Table 1. Note that for $J = 19$, $2^{19} = 524,228$ which corresponds to about 364

[5]Although there exist some one-minute data for ETH/USD and XRP/USD for earlier years, these data are replete with missing observations and hence are not reliable. For this reason, we begin the sample period in 2016 for ETH/USD and in 2018 for XRP/USD.

Figure 1: BTC/USD

The top panel displays the daily price and the bottom panel the daily volatility for BTC/USD. The sample period ranges from January 1, 2015 to October 15, 2020.

days. In other words, setting $J = 19$ effectively selects one year of data. For this reason, we set $J = 19$ and report results for each year of the sample (i.e., the "annual subsamples") for each currency. The exception is year 2020, when the data end on October 15, and therefore we have less than the required observations to implement $J = 19$. Consequently, for 2020 only, we set $J = 18$, which corresponds to $2^{18} = 262,144$ one-minute observations or about 182 days from the beginning of the year. Note that all annual subsamples are non-overlapping.

We begin our empirical analysis by reporting summary statistics for each of the three cryptocurrency exchange rates. Table 2 reports the results for BTC/USD, Table 3 for ETH/USD

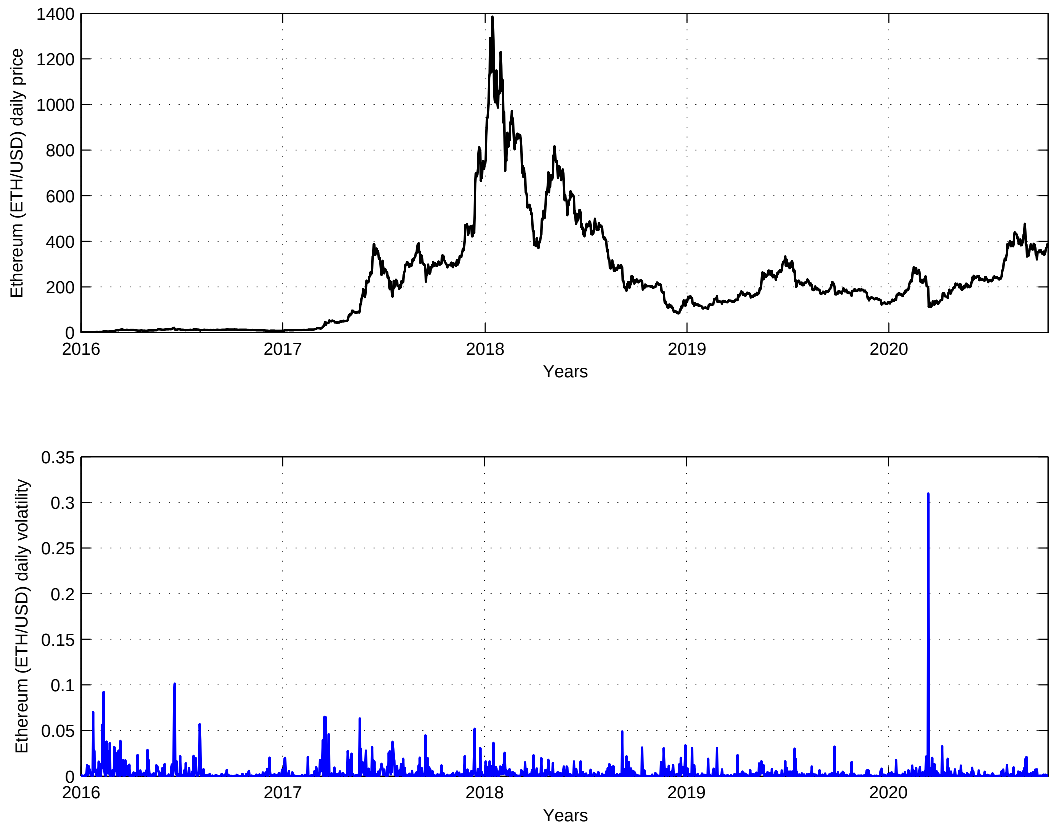

Figure 2: ETH/USD

The top panel displays the daily price and the bottom panel the daily volatility for ETH/USD. The sample period ranges from January 1, 2016 to October 15, 2020.

and Table 4 for XRP/USD. For these Tables, Panel A reports results for returns and Panel B for volatility. In each case, the Tables display summary statistics across four frequencies: 1-minute, 1-hour, 1-day, and 1-month. The mean and standard deviation are reported in annualized terms to facilitate comparison across frequencies. Moreover, in computing the statistics reported in Tables 2-4, we have winsorized the 1-minute returns at 1% to avoid the effect of excessive outliers at the highest frequency. Note that the returns of 1-hour, 1-day and 1-month are constructed by aggregating 1-minute returns.[6]

Our main findings can be summarized as follows. The mean annualized returns across

[6] See Brownlees and Gallo (2006) for further information on how to handle outliers in high-frequency data.

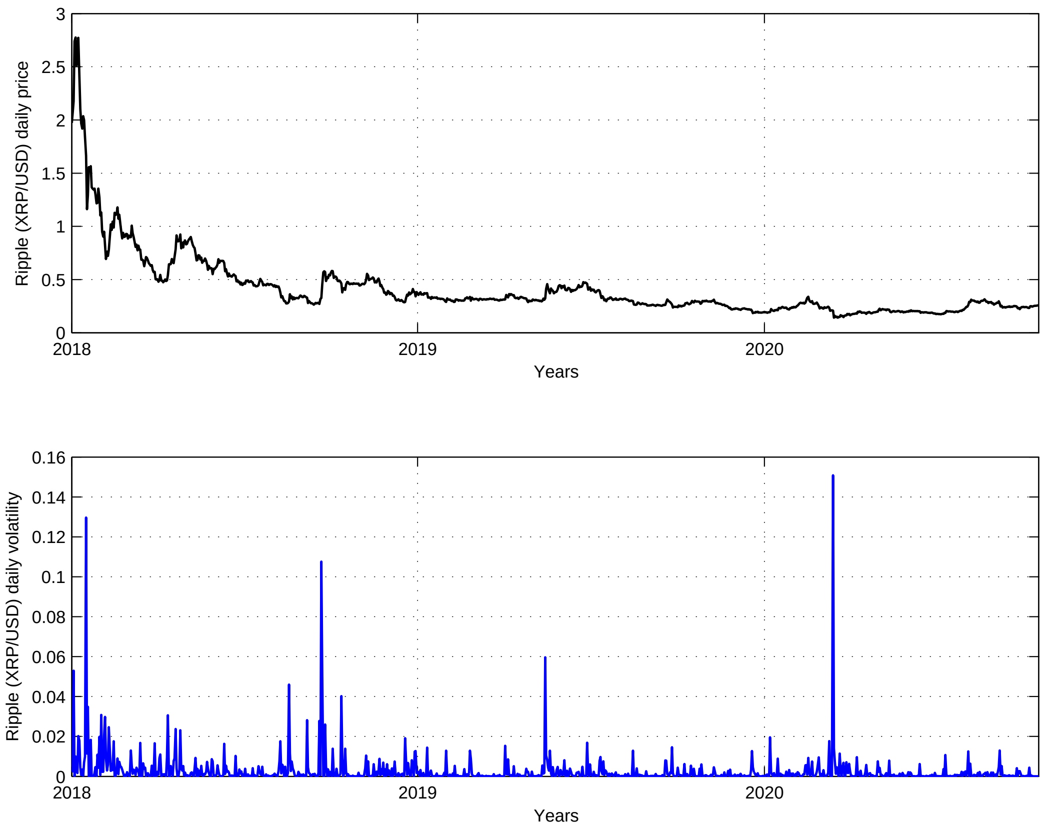

Figure 3: XRP/USD

The top panel displays the daily price and the bottom panel the daily volatility for XRP/USD. The sample period ranges from January 1, 2018 to October 15, 2020.

all cryptocurrencies are very large. For example, using annualized 1-month returns, the mean is 60.2% for BTC/USD, 123.4% for ETH/USD, and −76.6% for XRP/USD. The annualized standard deviation of returns is also very large: 72.9% for BTC/USD, 137.0% for ETH/USD, and 93.4% for XRP/USD. In addition, the higher the frequency, the more negative the skewness (i.e., the higher the skewness in absolute value), and the higher the kurtosis. For example, for ETH/USD, skewness is equal to −0.069 for 1-minute returns and 0.590 for 1-month returns, whereas kurtosis is equal to 8.588 for 1-minute returns and 3.315 for 1-month returns. In short, therefore, our results confirm for cryptocurrencies that the higher the frequency the further away from normal is the distribution of financial asset

Panel A: BTC/USD Returns				
	1-Minute	1-Hour	1-Day	1-Month
Mean (annualized)	1.164	0.724	0.717	0.602
St. Dev.(annualized)	0.524	0.651	0.693	0.729
Min	-0.003	-0.027	-0.115	-0.464
Max	0.003	0.025	0.111	0.575
Skewness	-0.137	-0.244	-0.188	0.103
Kurtosis	8.447	7.337	5.052	3.084
Panel B: BTC/USD Volatility				
Mean (annualized)	0.274	0.280	0.275	0.280
St. Dev. (annualized)	0.001	0.005	0.018	0.071
Min	0.000	0.000	2.3E-05	0.003
Max	8.4E-06	4.8E-04	0.008	0.092
Skewness	4.082	3.149	2.981	1.903
Kurtosis	20.078	15.623	14.936	6.530

Table 2: Summary Statistics: BTC/USD

The table presents summary statistics for BTC/USD returns and volatility for the following frequencies: 1-minute, 1-hour, 1-day and 1-month. The 1-minute returns are winsorized at 1%. The sample period ranges from January 1, 2015 to October 15, 2020.

returns.

For the volatility across all currencies, the results are similar. The main difference here is that skewness is consistently positive. As we move to a higher frequency both skewness and kurtosis become higher. For example, for ETH/USD, skewness rises from 0.607 for 1-month volatility to 4.076 for 1-minute volatility; and kurtosis rises from 2.276 for 1-month volatility to 19.994 for 1-minute volatility. In short, therefore, the choice of frequency is important in determining the statistical behaviour of cryptocurrency returns and volatility, especially for the higher-order moments of skewness and kurtosis. This motivates our analysis of volatility across different times scales. We turn to this analysis next.

4. Empirical Results

Our main empirical analysis focuses on implementing the wavelet HMT model to estimate the transition probabilities for two cases: (1) moving from a long horizon (low-frequency time scale) to a short horizon (high-frequency time scale), and (2) moving from a short horizon

Panel A: ETH/USD Returns				
	1-Minute	1-Hour	1-Day	1-Month
Mean (annualized)	1.236	1.375	1.394	1.234
St. Dev. (annualized)	0.775	1.056	1.066	1.370
Min	-0.004	-0.040	-0.161	-0.800
Max	0.004	0.041	0.181	1.115
Skewness	-0.069	0.028	0.269	0.590
Kurtosis	8.588	6.695	4.674	3.315
Panel B: ETH/USD Volatility				
Mean (annualized)	0.600	0.614	0.601	0.615
St. Dev. (annualized)	0.002	0.009	0.032	0.097
Min	0.000	0.000	8.5E-05	0.012
Max	1.8E-05	9.7E-04	0.015	0.011
Skewness	4.076	3.140	2.802	0.607
Kurtosis	19.994	16.709	15.105	2.276

Table 3: Summary Statistics: ETH/USD

The table presents summary statistics for ETH/USD returns and volatility for the following frequencies: 1-minute, 1-hour, 1-day and 1-month. The 1-minute returns are winsorized at 1%. The sample period ranges from January 1, 2016 to October 15, 2020.

(high-frequency time scale) to a long horizon (low-frequency time scale). Results for the two cases are reported separately. The transition probabilities associated with the second case are denoted as "reverse" transition probabilities to distinguish them from the transition probabilities of the first case.

The transition probabilities enable us to assess whether there is symmetric or asymmetric vertical dependence in cryptocurrency volatility. We define symmetric vertical dependence as the propagation of a given volatility state to the same volatility state (i.e., low-to-low and high-to-high) when moving from one time scale to the next. Formally, symmetric vertical dependence is the case when the low-to-low transition probability is significantly higher than the low-to-high transition probability *and* the high-to-high transition probability is significantly higher than the high-to-low transition probability. In other words, if volatility is low at one time scale and is expected to stay low at the next time scale and if volatility is high at one time scale and is expected to stay high at the next time scale, then vertical volatility dependence is symmetric.

If vertical volatility dependence is not symmetric, then it must be asymmetric. In other

Panel A: XRP/USD Returns				
	1-Minute	1-Hour	1-Day	1-Month
Mean (annualized)	-0.952	-1.074	-0.609	-0.766
St. Dev. (annualized)	0.791	0.912	0.908	0.934
Min	-0.004	-0.035	-0.141	-0.629
Max	0.004	0.036	0.162	0.564
Skewness	-0.049	0.017	0.149	0.501
Kurtosis	7.126	6.800	4.939	3.638
Panel B: XRP/USD Volatility				
Mean (annualized)	0.625	0.639	0.626	0.651
St. Dev. (annualized)	0.002	0.009	0.036	0.138
Min	0.000	0.000	0.001	0.013
Max	1.7E-05	0.001	0.017	0.208
Skewness	3.869	3.294	3.324	2.306
Kurtosis	18.391	16.904	17.715	9.748

Table 4: Summary Statistics: XRP/USD

The table presents summary statistics for XRP/USD returns and volatility for the following frequencies: 1-minute, 1-hour, 1-day and 1-month. The 1-minute returns are winsorized at 1%. The sample period ranges from January 1, 2018 to October 15, 2020.

words, all that is required for asymmetric volatility dependence is that either low-to-low is not significantly greater than low-to-high *or* high-to-high is not significantly greater than high-to-low. It is important to note that in all cases statistical significance is required to determine symmetric vertical dependence. Hence the absence of statistical significance in the conditions required for symmetric vertical dependence leads us to conclude that vertical dependence is asymmetric.

In what follows, we examine whether vertical volatility dependence is symmetric or asymmetric for two directions (long horizon to short horizon and vice versa) and three cryptocurrency exchange rates.

4.1. From Long to Short Horizons

4.1.1. Main Cryptocurrency: BTC/USD

We begin our discussion of the empirical results with evidence on volatility cascades from long to short horizons. This analysis addresses the question of whether there is a transmission

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2015	0.8571	0.1429	0.5300	0.4700
p-value	0.0000		0.4480	
2016	0.7772	0.2228	0.5885	0.4115
p-value	0.0000		0.0308	
2017	0.8535	0.1465	0.5517	0.4483
p-value	0.0000		0.2273	
2018	0.7936	0.2064	0.6478	0.3522
p-value	0.0000		0.0000	
2019	0.8103	0.1897	0.4895	0.5105
p-value	0.0000		0.7783	
2020	0.8347	0.1653	0.5750	0.4250
p-value	0.0000		0.0165	

Table 5: Average Transition Probabilities: BTC/USD

The table reports the average transition probabilities for BTC/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale j conditional on the volatility state in scale $j+1$. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

of volatility from long horizons (i.e., low-frequency trading) to short horizons (i.e., high-frequency trading). In other words, we assess whether volatility in economic fundamentals observed at a low frequency also affects volatility in high-frequency trading. To address this question, we estimate the transition probability of a volatility state at scale j given the volatility state at scale $j+1$.

The transition probabilities for BTC/USD are reported in Table 5. These transition probabilities are averaged across all time scales and are reported for annual non-overlapping subsamples. Recall that the BTC/USD data cover the 2015-2020 sample period. For the first five years (2015-2019), we use $J = 19$ scales, whereas for 2020 due to less data being available, we set $J = 18$. For each year, we report two-sided p-values of whether: (1) the low-to-low probability is significantly different from the low-to-high probability, and (2) the high-to-high probability is significantly different from the high-to-low probability.

Our main findings are as follows. For all years, the low-to-low transition probability is always significantly higher than the low-to-high probability. In other words, low volatility in economic fundamentals tends to be followed by low volatility in high-frequency trading with high likelihood and high statistical significance.

In contrast, the high-to-high evidence is mixed. For three of the six years (2016, 2018 and 2020), the high-to-high transition probability is greater than the high-to-low probability with high statistical significance. For these years, therefore, there is strong evidence of symmetric vertical dependence. For the remaining three years (2015, 2017 and 2019), although the high-to-high probability tends to be greater than 50%, it is not significantly higher than that high-to-low probability. Hence for these three years, the empirical evidence indicates asymmetric vertical dependence. In short, the empirical evidence indicates that there is symmetric vertical dependence for three years and asymmetric vertical dependence for another three years.

The results in Table 5 discussed above display transition probabilities, which are averaged across all time scales. To assess the transition probabilities across all scales (rather than on average), we focus on the 2020 results reported in Table 6. This table enriches our analysis by identifying for which scales we are more likely to observe a symmetric or asymmetric volatility cascade. The focus of Table 6 is on the year 2020 because this is the most recent year for which data are available and because the COVID-19 global pandemic provides a natural example of a low-frequency volatility shock. It would be interesting to see for which time scales high volatility due to COVID-19 was propagated to the higher frequency scales.[7]

The results in Table 6 indicate that for scale $j = 5$ and higher the low-to-low transitional probability is substantially greater than 50%. In fact, in most cases the low-to-low probability is close to 1. The high-to-high transitional probability is also greater than 50% for most (but not all) scales for $j = 7$ and higher. Hence the results in Table 6 strongly support that the vertical volatility dependence for BTC/USD in 2020 is symmetric across most scales.[8]

A concise way to visualize the volatility cascades from long to short horizons is with the graphs illustrated in Figure 4. This figure focuses on BTC/USD for the year 2020. The top

[7]Note that for Table 6 we report p-values only on the average transitional probabilities across scales as it is not possible to compute p-values for each time-scale separately.

[8]Note that for low scales (e.g., $j = 1$ to $j = 4$), the transitional probabilities are all 50%. This is natural since the lowest scales capture time horizons measured in a few minutes. For scales $j > 5$, the time horizon is measured in hours and the difference from one scale to the next becomes progressively larger.

Scale	Low-to-Low	Low-to-High	High-to-High	High-to-Low
18	0.9992	0.0008	0.4854	0.5146
17	0.9993	0.0007	0.6401	0.3599
16	0.9999	0.0001	0.1944	0.8056
15	0.9694	0.0306	0.5722	0.4278
14	0.9455	0.0545	0.9466	0.0534
13	0.9764	0.0236	0.4975	0.5025
12	0.9748	0.0252	0.5840	0.4160
11	0.9403	0.0597	0.7022	0.2978
10	0.9398	0.0602	0.8036	0.1964
9	0.9806	0.0194	0.7672	0.2328
8	0.9800	0.0200	0.8111	0.1889
7	0.9997	0.0003	0.6466	0.3534
6	0.7610	0.2390	0.2577	0.7423
5	0.5592	0.4408	0.4408	0.5592
4	0.5000	0.5000	0.5000	0.5000
3	0.5000	0.5000	0.5000	0.5000
2	0.5000	0.5000	0.5000	0.5000
1	0.5000	0.5000	0.5000	0.5000
mean	0.8347	0.1653	0.5750	0.4250
p-value	0.0000		0.0165	

Table 6: Scale-by-Scale Transition Probabilities: BTC/USD
The table reports the transition probabilities for BTC/USD volatility for the year 2020. The reported values are the probabilities of a given volatility state in scale j conditional on the volatility state in scale $j+1$. The p-values are for two-tailed t-tests of equality of means across all scales (between columns 2 and 3, then 4 and 5).

panel of Figure 4 plots volatility horizontally across time as well as vertically across time scales. Whenever the wavelet HMT model indicates that for a given time and time scale volatility is low, then volatility is plotted in white rectangles. When the model indicates volatility is high, it is plotted in black rectangles. Therefore, a heavy concentration of black areas in this graph indicates an outburst of volatility propagating vertically across the time scales. This is perhaps the best representation of a volatility cascade in the shape of a dark "waterfall" propagating across time scales. The bottom panel of the figure displays the daily volatility of BTC/USD across time. This allows us to associate the volatility cascades across time scales in the top panel with the standard volatility across time in the bottom panel.

Overall, Figure 4 provides strong support for symmetric vertical dependence. This is because black rectangles tend to be followed by black rectangles and white rectangles tend to be followed by white rectangles. Furthermore, we can observe that high volatility, which

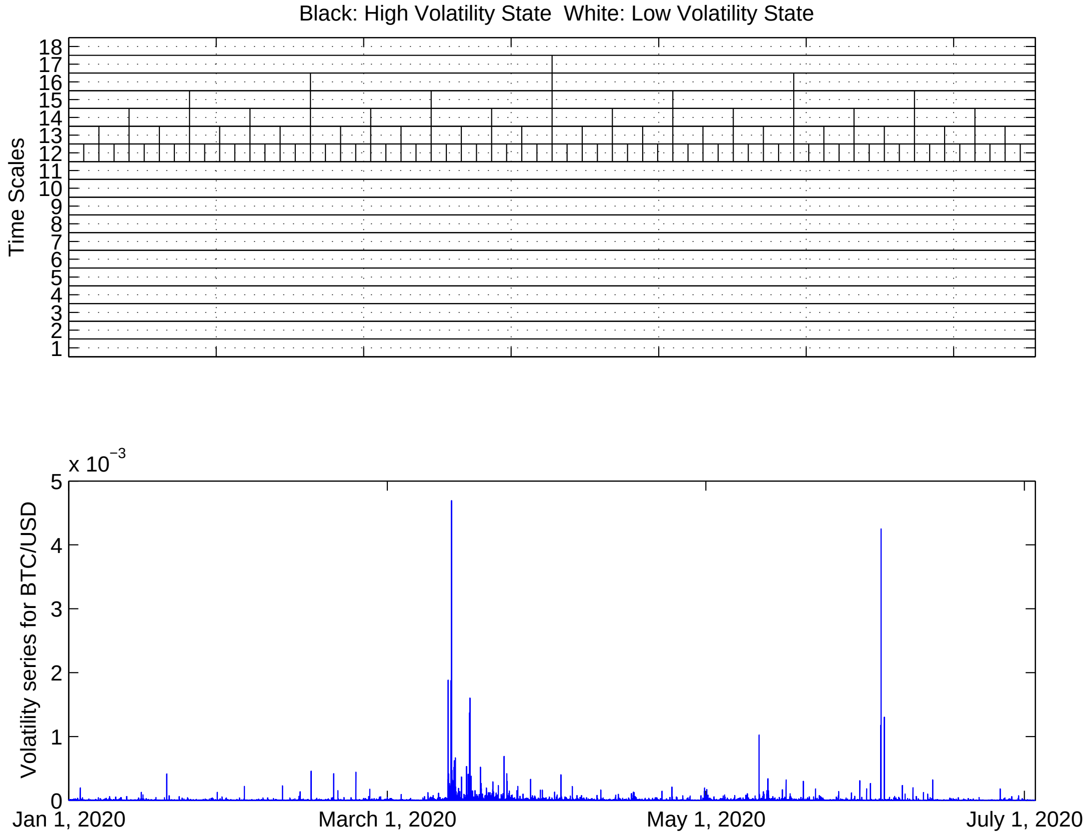

Figure 4: THE VOLATILITY CASCADE FOR BITCOIN IN 2020.

The top panel shows the volatility cascade for Bitcoin for the sample period of January 1, 2020 to July 3, 2020. This sample period corresponds to $J = 18$ time scales and 2^{18} one-minute returns. The white rectangles indicate a low-volatility state and the black rectangles indicate a high-volatility state. The bottom panel plots the volatility of one-minute Bitcoin returns across time.

typically begins around scale $j = 14$ (that corresponds to a horizon of about one month), tends to propagate to lower scales. This is a clear indication that volatility in economic fundamentals, which is typically observed on a monthly basis, is indeed associated with higher-frequency volatility. This leads to conclude that the vertical dependence in BTC/USD volatility is mostly symmetric.

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2016	0.8372	0.1628	0.6478	0.3522
p-value	0.0000		0.0000	
2017	0.8295	0.1705	0.4957	0.5043
p-value	0.0000		0.8999	
2018	0.8051	0.1949	0.6209	0.3791
p-value	0.0000		0.0026	
2019	0.7802	0.2198	0.5384	0.4616
p-value	0.0000		0.1902	
2020	0.8196	0.1804	0.5211	0.4789
p-value	0.0000		0.6196	

Table 7: Average Transition Probabilities: ETH/USD

The table reports the average transition probabilities for ETH/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale j conditional on the volatility state in scale $j+1$. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2018	0.8050	0.1950	0.5869	0.4131
p-value	0.0000		0.0190	
2019	0.8114	0.1886	0.5784	0.4216
p-value	0.0000		0.0090	
2020	0.8065	0.1935	0.5199	0.4801
p-value	0.0000		0.6270	

Table 8: Average Transition Probabilities: XRP/USD

The table reports the average transition probabilities for XRP/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale j conditional on the volatility state in scale $j+1$. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

4.1.2. Other Cryptocurrencies: ETH/USD and XRP/USD

Next we turn to the empirical results related to the other two cryptocurrencies. Table 7 reports the results for ETH/USD for the years 2016-2020, and Table 8 is for XRP/USD for the years 2018-2020.

In short, the results for ETH/USD and XRP/USD support the results for BTC/USD in two ways. First, in all cases the low-to-low transitional probability is statistically significantly greater than the low-to-high probability. Second, the high-to-high transitional probability is often significantly greater than the high-to-low probability. This is the case for 2016 and

2018 for ETH/USD and for 2018 and 2019 for XRP/USD. For the remaining years, the high-to-high probability tends to be more than 50%, but it is not statistically significantly different from the high-to-low probability.

It is interesting to note that across all cryptocurrencies, the years 2016 and 2018 seem to display the strongest evidence of a symmetric volatility cascade. The year 2016 was marked by pronounced volatility in the Chinese stock market and a heightened interest in cryptocurrencies whose prices rose to unprecedented levels. This was followed by a market crash in late 2017 and early 2018. In brief, therefore, when low-frequency volatility was high, the symmetric volatility cascade was pronounced across all cryptocurrencies. A notable exception is the COVID-19 pandemic in 2020, which caused a symmetric volatility cascade for BTC/USD, but not for the other two cryptocurrencies.

4.1.3. Concluding Remarks

In conclusion, the empirical results for volatility transmission from long to short horizons indicate that vertical dependence tends to be symmetric, especially for years of high volatility. This is a new result that is specific to cryptocurrencies. For instance, this evidence is contrary to Gençay *et al.* (2010), who show that for standard exchange rates (e.g., the US dollar - German mark exchange rate) the vertical volatility dependence is exclusively asymmetric. For this reason, we make a novel contribution to the literature.

4.2. From Short to Long Horizon

In this section, we examine volatility cascades in the reverse direction: from short to long horizons. This analysis addresses the question of whether there is a transmission of volatility from short horizons (i.e., high-frequency trading) to long horizons (i.e., low-frequency trading). The bi-directional approach to the study of volatility cascades is also a novel contribution to this body of research since the literature typically considers only one-directional volatility cascades from long to short horizons similar to what this paper has done in the

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2015	0.7680	0.2320	0.1926	0.8074
p-value	0.0000		0.0000	
2016	0.7743	0.2257	0.2212	0.7788
p-value	0.0000		0.0000	
2017	0.7344	0.2656	0.2327	0.7673
p-value	0.0000		0.0000	
2018	0.7880	0.2120	0.2121	0.7879
p-value	0.0000		0.0000	
2019	0.7613	0.2387	0.1879	0.8121
p-value	0.0000		0.0000	
2020	0.8029	0.1971	0.1984	0.8016
p-value	0.0000		0.0000	

Table 9: Reverse Transition Probabilities: BTC/USD

The table reports the reverse average transition probabilities for BTC/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale $j+1$ conditional on the volatility state in scale j. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

previous section.

Our empirical analysis is based on the reverse average transition probabilities. These are reported in Table 9 for BTC/USD. Furthermore, Table 10 provides details on the scale-by-scale reverse transition probabilities for BTC/USD for the year 2020. The reverse average transition probabilities for ETH/USD are reported in Table 11, and for XRP/USD in Table 12.

All reverse transition probabilities across all cryptocurrencies lead to the same findings and hence we discuss them all together. Specifically, there are two important findings in this analysis. First, the low-to-low reverse average transition probabilities are always substantially higher than 50%. Also the difference between the low-to-low probability and the low-to-high probability is always statistically significant. Therefore, this evidence overwhelmingly indicates that low volatility at a high-frequency time scale is very likely to lead to low volatility at a lower frequency time scale.

Second, the high-to-low reverse average transition probabilities are also substantially higher than 50% for all cryptocurrencies. The difference between the high-to-low probability and the high-to-high probability is statistically significant in all cases. Therefore, the em-

Scale	Low-to-Low	Low-to-High	High-to-High	High-to-Low
18	0.9897	0.0103	0.0072	0.9928
17	0.9941	0.0059	0.0001	0.9999
16	0.9869	0.0131	0.0097	0.9903
15	0.9709	0.0291	0.0244	0.9756
14	0.9958	0.0042	0.0001	0.9999
13	0.9368	0.0632	0.0765	0.9235
12	0.8983	0.1017	0.1457	0.8543
11	0.9372	0.0628	0.0215	0.9785
10	0.7824	0.2176	0.2793	0.7207
9	0.8564	0.1436	0.4594	0.5406
8	0.7801	0.2199	0.2864	0.7136
7	0.7083	0.2917	0.0000	1.0000
6	0.8713	0.1287	0.0010	0.9990
5	0.7436	0.2564	0.2616	0.7384
4	0.5004	0.4996	0.4995	0.5005
3	0.5000	0.5000	0.5000	0.5000
2	0.5000	0.5000	0.5000	0.5000
1	0.5000	0.5000	0.5000	0.5000
mean	0.8029	0.1971	0.1984	0.8016
p-value	0.0000		0.0000	

Table 10: Scale-by-Scale Reverse Transition Probabilities: BTC/USD

The table reports the reverse transition probabilities for BTC/USD volatility for the year 2020. The reported values are the probabilities of a given volatility state in scale $j + 1$ conditional on the volatility state in scale j. The p-values are for two-tailed t-tests of equality of means across all scales (between columns 2 and 3, then 4 and 5).

pirical evidence overwhelmingly indicates that high volatility at short horizons is *not* likely to propagate to high volatility associated with long horizons. In other words, high volatility associated with high-frequency trading is unlikely to lead to high volatility for low-frequency traders.

Taken together, the two findings indicate that there is strong asymmetry in reverse volatility cascades. This asymmetric vertical dependence in cryptocurrency volatility is consistent with the Heterogeneous Market Hypothesis of Müller *et al.* (1997), which proposes that high short-term volatility does not cause high long-term volatility. Our findings lend strong support to the Heterogeneous Market Hypothesis. These findings are robust across all cryptocurrencies and all annual subsamples, including the 2017-2018 Bitcoin market crash and the 2020 COVID-19 global pandemic.

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2016	0.8313	0.1687	0.2804	0.7196
p-value	0.0000		0.0000	
2017	0.7109	0.2891	0.2495	0.7505
p-value	0.0000		0.0000	
2018	0.8082	0.1918	0.1985	0.8015
p-value	0.0000		0.0000	
2019	0.7926	0.2074	0.2356	0.7644
p-value	0.0000		0.0000	
2020	0.8321	0.1679	0.2721	0.7279
p-value	0.0000		0.0000	

Table 11: Reverse Transition Probabilities: ETH/USD

The table reports the reverse average transition probabilities for ETH/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale $j+1$ conditional on the volatility state in scale j. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

Year	Low-to-Low	Low-to-High	High-to-High	High-to-Low
2018	0.7878	0.2122	0.2584	0.7416
p-value	0.0000		0.0000	
2019	0.7699	0.2301	0.2536	0.7464
p-value	0.0000		0.0000	
2020	0.8229	0.1771	0.1668	0.8332
p-value	0.0000		0.0000	

Table 12: Reverse Transition Probabilities: XRP/USD

The table reports the reverse average transition probabilities for XRP/USD volatility. The reported values are the average across all time scales of the probability of a given volatility state in scale $j+1$ conditional on the volatility state in scale j. The p-values are for two-tailed t-tests of equality of means (between columns 2 and 3, then 4 and 5).

5. Conclusion

This paper studies how volatility in three popular cryptocurrencies cascades through different investor horizons. In other words, instead of the usual approach of studying the dynamics of volatility horizontally across time, our analysis studies the transmission of volatility vertically across different time scales. This is a new approach to the study of cryptocurrency volatility. It is based on a rich data set of one-minute returns for Bitcoin, Ethereum and Ripple against the US dollar. These data allow us to consider several time scales ranging from one minute to about one year. The approach is also based on a well-established methodology that implements the wavelet HMT model to estimate the transition probability of low or high volatility from one scale to the next. In essence, the question we address is the following: does high volatility at long horizons lead to high volatility at short horizons? And does the answer to this question change if we move in the reverse direction, i.e., from short to long horizons? Our empirical analysis is designed to answer these questions.

We show that volatility clustering represents an important stylized fact about cryptocurrency volatility, not only in the time domain, but also in the frequency domain. In particular, we find that for the direction of long horizons to short horizons, the volatility cascade tends to be mostly symmetric. This implies that low volatility at long horizons is likely to be followed by low volatility at short horizons and high volatility at long horizons is likely to be followed by high volatility at short horizons. This symmetric vertical dependence in volatility is a new result that is specific to cryptocurrencies. For instance, when using standard exchange rate data, the empirical literature has established an asymmetric vertical dependence.

Our analysis is bi-directional in that we also examine the vertical transmission of volatility in the reverse direction: from short to long horizons. In this case, we find that for cryptocurrencies there is strong evidence of asymmetric vertical dependence. In other words, low volatility at short horizons is still likely to be followed by low volatility at long horizons, but high volatility at short horizons is likely to be followed by low volatility at long horizons. This is an important result because it implies that an outburst of high-frequency volatility is not likely to cause a similar outburst of low-frequency volatility.

Taken all together, our findings indicate that low volatility is likely to persist across time scales no matter what direction we take. In contrast, high volatility at long horizons persists when moving to short horizons, but high volatility at short horizons does not persist when moving to long horizons. These findings are consistent with the Heterogeneous Market Hypothesis of Müller *et al.* (1997). The economic mechanism supporting this hypothesis asserts that high long-term volatility may cause high short-term volatility but not vice versa. Based on one-minute data, our analysis confirms that this is indeed the case for cryptocurrencies and, therefore, it strongly supports using a vertical structure that discriminates across short and long horizons in the study of volatility.

References

Adcock, R. and Gradojevic, N. (2019). Non-fundamental, non-parametric Bitcoin forecasting. *Physica A*, **531**, 121727.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, **96**, 42–55.

Ardia, D., Bluteau, K., and Rüede, M. (2019). Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters*, **29**, 266–271.

Auer, R. and Claessens, S. (2018). Regulating cryptocurrencies: Assessing market reactions. *BIS Quarterly Review*, **September 2018**, 51–65.

Bacry, E., Kozhemyak, A., and Muzy, J.-F. (2008). Continuous cascade models for asset returns. *Journal of Economic Dynamics and Control*, **32**, 156–199.

Bhambhwani, S. M., Delikouras, S., and Korniotis, G. M. (2020). Blockchain characteristics and the cross-section of cryptocurrency returns. SSRN working paper 3342842.

Borri, N. (2019). Conditional tail-risk in cryptocurrency markets. *Journal of Empirical Finance*, **50**, 1–19.

Borri, N. and Shakhnov, K. (2020a). The cross-section of cryptocurrency returns. SSRN working paper 3241485.

Borri, N. and Shakhnov, K. (2020b). Regulation spillovers across cryptocurrency markets. *Finance Research Letters*, **36**, 101333.

Bouri, E., Lau, C. K. M., Lucey, B., and Roubaud, D. (2019). Trading volume and the predictability of return and volatility in the cryptocurrency market. *Finance Research Letters*, **29**, 340–346.

Bouri, E., Roubaud, D., and Shahzad, S. J. H. (2020). Do Bitcoin and other cryptocurrencies jump together? *Quarterly Review of Economics and Finance*, **76**, 396–409.

Briere, M., Oosterlinck, K., and Szafarz, A. (2015). Virtual currency, tangible return: Portfolio diversification with Bitcoin. *Journal of Asset Management*, **16**, 365–373.

Brownlees, C. and Gallo, G. (2006). Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational Statistics and Data Analysis*, **51**, 2232–2245.

Calvet, L. and Fisher, A. (2002). Multifractality in asset returns: Theory and evidence. *Review of Economics and Statistics*, **84**, 381–406.

Celeste, V., Corbet, S., and Gurdgiev, C. (2020). Fractal dynamics and wavelet analysis: Deep volatility and return properties of Bitcoin, Ethereum and Ripple. *Quarterly Review of Economics and Finance*, **76**, 310–324.

Charfeddine, L., Benlagha, N., and Maouchi, Y. (2020). Investigating the dynamic relationship between cryptocurrencies and conventional assets: Implications for financial investors. *Economic Modelling*, **85**, 198–217.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, **7**, 174–196.

Crouse, M. S., Nowak, R. D., and Baraniuk, R. G. (1998). Wavelet-based statistical signal processing using hidden Markov models. *IEEE Transactions on Signal Processing*, **46**, 886–902.

Dyhrberg, A. H. (2016). Bitcoin, gold and the dollar − A GARCH volatility analysis. *Finance Research Letters*, **16**, 85–92.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*, **134**, 91–109.

Eisl, A., Gasser, S. M., and Weinmayer, K. (2020). Caveat emptor: Does Bitcoin improve portfolio diversification? SSRN working paper 2408997.

Gençay, R., Gradojevic, N., Selçuk, F., and Whitcher, B. (2010). Asymmetry of information flow between volatilities across time scales. *Quantitative Finance*, **10**, 895–915.

Gronwald, M. (2019). Is Bitcoin a commodity? On price jumps, demand shocks, and certainty of supply. *Journal of International Money and Finance*, **97**, 86–92.

Hafner, C. M. (2020). Testing for bubbles in cryptocurrencies with time-varying volatility. *Journal of Financial Econometrics*, **18**, 233–249.

Hu, A. S., Parlour, C. A., and Rajan, U. (2019). Cryptocurrencies: Stylized facts on a new investible instrument. *Financial Management*, **48**, 1049–1068.

Kang, S. H., McIver, R., and Hernandez, J. A. (2019). Co-movements between Bitcoin and Gold: A wavelet coherence analysis. *Physica A*, **536**, 120888.

Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, **158**, 3–6.

Klein, T., Thu, H. P., and Walther, T. (2018). Bitcoin is not the new gold - A comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis*, **59**, 105–116.

Kristoufek, L. (2015). What are the main drivers of the Bitcoin price? Evidence from wavelet coherence analysis. *PLoS One*, **10**, e0123923.

Liu, W., Liang, X., and Cui, G. (2020). Common risk factors in the returns on cryptocurrencies. *Economic Modelling*, **86**, 299–305.

Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, **135**, 293–319.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, **36**, 394–419.

Mandelbrot, B. (1997). *Fractals and Scaling in Finance*. Springer, New York.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., and von Weizsäcker, J. E. (1997). Volatilities of different time resolutions - Analyzing the dynamics of market components. *Journal of Empirical Finance*, **4**, 213–239.

Panagiotidis, T., Stengos, T., and Vravosinos, O. (2018). On the determinants of Bitcoin returns: A LASSO approach. *Finance Research Letters*, **27**, 235–240.

Scaillet, O., Treccani, A., and Trevisan, C. (2020). High-frequency jump analysis of the Bitcoin market. *Journal of Financial Econometrics*, **18**, 209–232.

Segnon, M. and Bekiros, S. (2020). Forecasting volatility in Bitcoin market. *Annals of Finance*, **16**, 435–462.

Trimborn, S. and Hardle, W. K. (2018). CRIX an index for cryptocurrencies. *Journal of Empirical Finance*, **49**, 107–122.

Trimborn, S., Li, M., and Hardle, W. K. (2020). Investing with cryptocurrencies - A liquidity constrained investment approach. *Journal of Financial Econometrics*, **18**, 280–306.