Article

Bitcoin at High Frequency

Leopoldo Catania [1,]* and Mads Sandholdt [2]

[1] Department of Economics and Business Economics, Aarhus University and CREATES, Aarhus BSS,
   Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark
[2] Tvilum A/S, Egon Kristiansens Allé 2, DK-8882 Faarvang, Denmark; mads.sandholdt@me.com
* Correspondence: leopoldo.catania@econ.au.dk; Tel.: +44-8716-5536

check for
updates

Abstract: This paper studies the behaviour of Bitcoin returns at different sample frequencies. We consider high frequency returns starting from tick-by-tick price changes traded at the Bitstamp and Coinbase exchanges. We find evidence of a smooth intra-daily seasonality pattern, and an abnormal trade- and volatility intensity at Thursdays and Fridays. We find no predictability for Bitcoin returns at or above one day, though, we find predictability for sample frequencies up to 6 h. Predictability of Bitcoin returns is also found to be time–varying. We also study the behaviour of the realized volatility of Bitcoin. We document a remarkable high percentage of jumps above 80%. We also find that realized volatility exhibits: (i) long memory; (ii) leverage effect; and (iii) no impact from lagged jumps. A forecast study shows that: (i) Bitcoin volatility has become more easy to predict after 2017; (ii) including a leverage component helps in volatility prediction; and (iii) prediction accuracy depends on the length of the forecast horizon.

Keywords: bitcoin; realized volatility; HAR; high frequency

1. Introduction

One of the reasons why cryptocurrencies—and in particular Bitcoin introduced by Nakamoto (2009)—became so popular in 2017 has been their huge price increase which caught the attention from both the media and regular people. Indeed, we find that approximately 65% of all Bitcoin transactions happened in 2017 or later.[1] As a consequence of this huge interest, Bitcoin experienced a price increase of 1324% from the begin to the end of 2017. The financial industry and the academics have also been very interested in Bitcoin over the last years. For example, the Chicago Mercantile Exchange (CME) as well as Nasdaq and the Tokyo Financial Exchange started to negotiate Bitcoin futures during 2017 and 2018, see CME (2017), Bloomberg (2017), and Cryptocoinsnews (2017). Academics working in the field of financial econometrics have studied Bitcoin using well known methodologies, such as ARMA–GARCH models. For example, Dyhrberg (2016) compared Bitcoin with gold and the American dollar and classified the behavior of Bitcoin in between these two assets. Bariviera (2017) found long memory in the Bitcoin volatility measured as the logarithmic difference between intraday highest and lowest prices. Phillip et al. (2018) document long memory in cryptocurrencies as well. Additional results about the time–dependence properties of cryptocurrencies are reported in Zhang et al. (2019). Ardia et al. (2018) and Stavroyiannis (2018) model and forecast the value at risk for Bitcoin. Recently, Catania and Grassi (2017) show that standard volatility models, like GARCH, are generally not suitable for cryptocurrency time–series and suggest to use a more sophisticated modelling technique based on the score driven approach, see Creal et al. (2013) and Harvey (2013). The predictability of

---

[1] Specifically, 54.42% for the Bitstamp exchange and 68.78% for the Coinbase exchange.

cryptocurrencies returns and volatility has been studied in Catania et al. (2019) and Catania et al. (2018), respectively. Understanding the behavior of Bitcoin volatility has important implications for individual investors and public institutions. Individual investors – people not in the finance industry but interested in trading cryptocurrencies – should be informed about how the Bitcoin volatility evolves and whether these investment opportunities match with their risk profile. Public institutions, like the central banks of Ecuador, Tunisia, and Sweden who are considering issuing their own cryptocurrency, should be interested in the behavior of Bitcoin due to the possible systemic risk they would face by entering in this new market.

In this paper, we study the Bitcoin returns at high frequency and its realized volatility measure. We start by describing the construction of our dataset from the raw transactions downloaded from the Bitstamp and Coinbase exchanges.[2] The first part of our analysis focuses on the in sample properties of Bitcoin returns sampled at different frequencies. We study the autocorrelation structure of Bitcoin returns as well as the intraday and intraweek seasonalities of Bitcoin volatility and traded volumes. Results indicate strong presence of both intraday and intraweek seasonality in the volatility. We also document that seasonality is different across exchanges. Specifically, Coinbase follows the US trading activity, while Bitstamp the European one. The second part the paper focuses on predicting realized volatility for Bitcoin using several forecasting models. Realized volatility is a consistent estimator of the quadratic variation of the price process and is presently widely used in financial and risk management applications, see Bauwens et al. (2012) for a recent overview. We consider the baseline HAR-RV model of Corsi (2008) as well as its generalization with the inclusion of jumps (Andersen et al. 2007; Barndorff-Nielsen 2004; Barndorff-Nielsen and Shephard 2003; Barndorff-Nielsen et al. 2006) and the leverage component (Corsi et al. 2012). Our results suggest that: (i) the predictability of Bitcoin realized volatility has increased over time; (ii) including the leverage component helps in predicting future volatility levels; and (iii) predictability varies with the forecast horizon. Both in sample and out of sample results are reported for the two exchanges as well as for the subperiod 2017–2018 which coincides with the explosion in the interest of Bitcoin. Overall, our results indicate several peculiarities of Bitcoin volatility compared to the volatility of alternative investment opportunities. First, we find that the frequency of jumps is much higher compared to common findings. Second, in contrast to Catania and Grassi (2017) and Ardia et al. (2018) who find an "inverted" leverage effect, our results show that Bitcoin exhibits a leverage effect similar to that of equity assets when this is measured using the Realized Volatility estimator. This last point also supports the arguments of Dyhrberg (2016) who classify Bitcoin as an asset and not as an exchange rate. Indeed, the leverage effect is of little importance for exchange rates as documented for example by Hansen and Lunde (2005) and Ardia et al. (2018).

The structure of the paper is organized as follows. Section 2 details the dataset we build starting from tick–by–tick price changes to equally spaced logarithmic returns. Section 3 studies the behaviour of Bitcoin returns sampled at different frequencies. Section 4 analyses the realized volatility of Bitcoin. Conclusions are drawn in Section 5.

2. Data

Our dataset is composed by tick-by-tick traded prices recorded at the two exchanges Bitstamp and Coinbase over which the majority of the transactions takes place. From both exchanges we collect the tick-by-tick transaction data using the freely available API at www.api.bitcoincharts.com. The raw data include the transaction price for every trade, along with the amount of Bitcoins traded. Observations

---

[2]   Among several exchanges where Bitcoin is traded we have selected two of the most active ones. Another possibility would have been to consider GDAX instead of Coinbase since this exchange might give a better representation of Bitcoin. However, we have decided to use Coinbase since, differently from GDAX and Bitstamp which are traditional exchanges, it is a click and buy exchange which allows investor to immediately invest in Bitcoin. We thank an anonymous referee for pointing out this to us.

start the 13 September 2011 on Bitstamp and the 1 December 2014 on Coinbase and are reported in UTC time. We record data up to 18 March 2018 for a total of 22,457,894 trades for Bitstamp and 39,439,004 for Coinbase. Unfortunately, trades are reported with a precision of one second. However, even though it is not reported in the exchanges documentations, we conjecture that trades within a second are reported in a chronological order. Overall, we find that 47.35% and 70.40% of the recorded transactions happen simultaneously with at least one other trade for Bitstamp and Coinbase, respectively.

Data Cleaning

　　Similar to standard high-frequency financial time-series, raw tick-by-tick Bitcoin prices are contaminated by wrongly reported observations. As suggested by Barndorff-Nielsen et al. (2009), we start by removing all those transactions associated with zero or negative volume. As a second step of data cleaning, we apply the methodology of Brownlees and Gallo (2006) to filter each transaction price. Specifically, let p_i be the price of Bitcoin associated with trade i in our dataset, we apply the following rule:

$$(|p_i - \bar{p}_i(k)| < 3s_i(k) + \gamma) = \begin{cases} \text{True,} & p_i \text{ is kept} \\ \text{False,} & p_i \text{ is removed} \end{cases}, \tag{1}$$

where $\bar{p}_i(k)$ and $s_i(k)$ denote δ-trimmed mean and sample standard deviation of a neighborhood of k observations around i, respectively. According to Brownlees and Gallo (2006), the positive integer k should be chosen as a function of the trading intensity, which for Bitcoin is relative high since there is no minimum buy-level. The additional tuning parameter $\gamma \in (0,1)$ is a granularity parameter and should prevent a zero standard deviation caused by sequences of k equal prices. Finally, $\delta \in (0,1)$ helps to reduce the effect of extreme observations during the filtering procedure. We run the filter reported in Equation (1) using different choices of γ, and k. Table 1 reports results from the cleaning procedure in terms of percentage of outliers eliminated over the full sample and on a daily basis, for different choices of the tuning parameters $\gamma \in \{0.02, 0.04, 0.06\}$ and $k \in \{40, 60, 80\}$. We set $\delta = 5\%$ and found that results are robust to this choice.

Table 1. Number of outliers and average number of outliers per day in percentage points as a function of the rolling trimmed mean parameter k and the granularity parameter γ. Results are reported for Coinbase and Bitstamp using Equation (1) with $\delta = 5\%$.

(k, γ)	$(40, 0.02)$	$(40, 0.04)$	$(40, 0.06)$	$(60, 0.02)$	$(60, 0.04)$	$(60, 0.06)$	$(80, 0.02)$	$(80, 0.04)$	$(80, 0.06)$
				Bitstamp					
Number of outliers	141,029	134,584	128,934	104,018	99,007	85,593	82,623	78,603	75,438
Average outliers per day	0.51%	0.46%	0.42%	0.38%	0.34%	0.30%	0.31%	0.27%	0.25%
				Coinbase					
Number of outliers	118,162	107,936	100,570	98,037	89,922	84,009	85,499	78,482	73,521
Average outliers per day	0.24%	0.20%	0.17%	0.19%	0.16%	0.14%	0.17%	0.14%	0.12%

　　The percentage of outliers we find are consistent with the findings of Brownlees and Gallo (2006) and Barndorff-Nielsen et al. (2009). For $k = 80$ and $\gamma = 0.06$ we find almost half the amount of outliers than for the case $k = 40$ and $\gamma = 0.02$. In general, we see a clear decreasing pattern in the amount of outliers when increasing the surrounding neighborhood as well as the granularity parameter. Consistently with Brownlees and Gallo (2006), we also find that the amount of outliers increases along with the increase in the number transactions. Indeed, for the case $k = 60$ and $\gamma = 0.02$, we find that 58.66% and 81.95% of the outliers are from 2017 or later for Bitstamp and Coinbase, respectively. Generally, we find that the outcome of the cleaning procedure is similar to that of Brownlees and Gallo (2006) suggesting that no particular attention should be made when dealing with high-frequency Bitcoin prices compared to the standard procedure employed for other financial series. Therefore, we stick to Brownlees and Gallo (2006) and use price series filtered using $(k, \gamma) = (60, 0.02)$ for our analysis. Starting from the filtered series of prices, we compute an equally spaced sequence of one–second prices and volumes. When multiple transactions are available within the same second, we

set the final price to the median price computed over that second. Days with less then 40 observations have been removed for the dataset. We removed 167 days for Bitstamp and 16 days for Coinbase from the begin of the sample. Beside this, in order to remain with a time series without missing days, we let our data set to start from 17 March 2013 for Bitstamp and from 2 February 2015 for Coinbase. In order to isolate the effect of 2017, for the rest of the paper results are reported for the full sample as well as for the sub-sample after 1 January 2017 at midnight labelled as "Hype". Table 2 summarizes the results from the cleaning procedure.

Table 2. This table summarizes the results from the cleaning procedure. Starting from the raw dataset, transactions associated with non positive volume are removed. Outliers are identified following the procedure of Brownlees and Gallo (2006). The "Simultaneous ticks" reports the number of transactions that occurs within the same second. The row "Final sample size in seconds" and " Trading days" report the number of equally spaced observations in second and the number of trading days, respectively. Results are reported for the two exchanges Bitstamp and Coinbase as well as for the full sample and the sub–sample "Hype".

	Bitstamp		Coinbase	
	Full Sample	**Hype**	**Full Sample**	**Hype**
Raw observations	22,346,195	12,217,195	39,285,138	27,126,897
Volume ≤ 0	1112	98	0	0
Outliers	104,295	61,350	103,071	85,418
Simultaneous ticks	9,976,748	5,338,296	28,400,400	18,861,682
Final sample size in seconds	12,263,584	6,817,445	10,781,667	8,179,797
Trading days	1825	442	1132	442

3. High Frequency Bitcoin Returns and Realized Volatiltiy

We start our analysis by computing the series of percentage Bitcoin logarithm returns at second $s = 1, \ldots, 86,400$ in day t as:

$$r_{s,t} = 100 \times [\log(p_{s,t}) - \log(p_{s-1,t})],$$

where $p_{s,t}$ and $p_{s-1,t}$ are two subsequent Bitcoin prices for day t. One second logarithmic returns are subsequently aggregated at different frequencies as reported below. We also compute the realized volatility for day t as in Andersen et al. (2001b). As suggested by Liu et al. (2015) realized volatility is computed using 5-min returns, as:

$$RV_t = \sum_{j=1}^{288} \tilde{r}_{j(300),t}^2,$$

where $\tilde{r}_{j(N),t} = \sum_{s=\underline{m}_j}^{\overline{m}_j} r_{s,t}$ and $\underline{m}_j = (j-1) \times N + 1$ and $\overline{m}_j = j \times N$ and we set $N = 300$ to achieve 5-min. aggregation. Table 3 reports a comparison between the realized variance of Bitcoin and the variance of S&P 500 measured with the VIX. We note that the volatility of Bitcoin and that of S&P 500 are comparable during the years 2015 and 2016. On the contrary, in the period 2017–2018 the volatility of Bitcoin is considerably higher compared to that of S&P 500. We also note that volatility of Bitcoin is higher during the bear market period of 2018 than during the bubble period of 2017. Overall, results indicate that the volatility of volatility is much higher for Bitcoin than for the S&P 500.

Descriptive statistics for Bitcoin percentage log returns aggregated at the 5-min. frequency are reported in Table 4. We observe that both series are characterized by extreme observations. We find that returns traded at the Bitstamp exchange exhibit higher volatility. However, we also note that returns traded at Coinbase are characterized by more pronounced negative skewness and higher excess of kurtosis. Overall, departure from Gaussianity is evident from the data. To conclude the analysis of Bitcoin returns at 5-min, in Figure 1 we report Gaussian kernel densities estimated on the mean,

standard deviation, skewness, and kurtosis coefficients computed over each trading day available in our dataset. Interestingly, while the skewness and the excess of kurtosis coefficients are similar across exchanges and sub–samples, we note that the distribution of the standard deviation is considerably shifted to the right during the Hype period. Furthermore, it is also evident that during the Hype period we observe more dispersed average returns.

Table 3. Comparison over the period 2015–2018 between daily average realized variance of Bitcoin, measured using Coinbase and Bitstamp 5-min returns and S&P 500 measured using the volatility index VIX. The first four columns report the sample while the last four columns the standard deviation (i.e., the volatility of volatility).

	Mean				Standard Deviation			
	2015	2016	2017	2018	2015	2016	2017	2018
S&P 500	16.67	15.83	11.09	16.64	4.34	3.97	1.36	5.09
Coinbase	11.25	6.89	37.24	59.02	29.76	17.26	78.46	67.15
Bitstamp	23.10	10.08	37.46	64.88	72.48	14.69	57.19	62.44
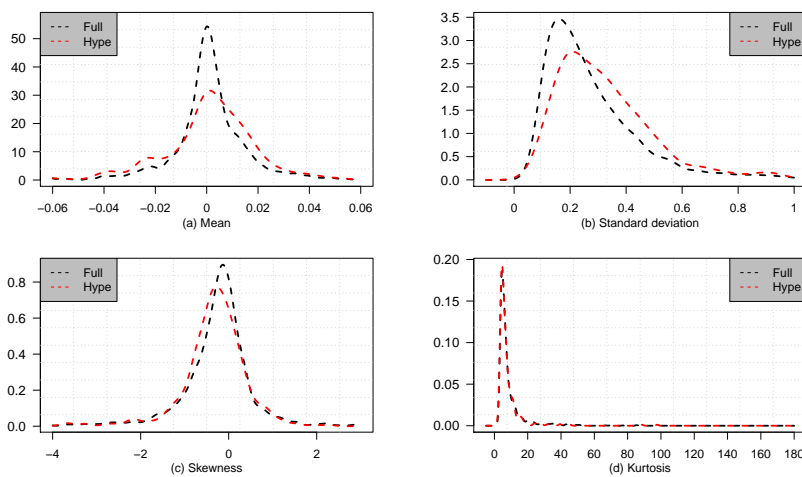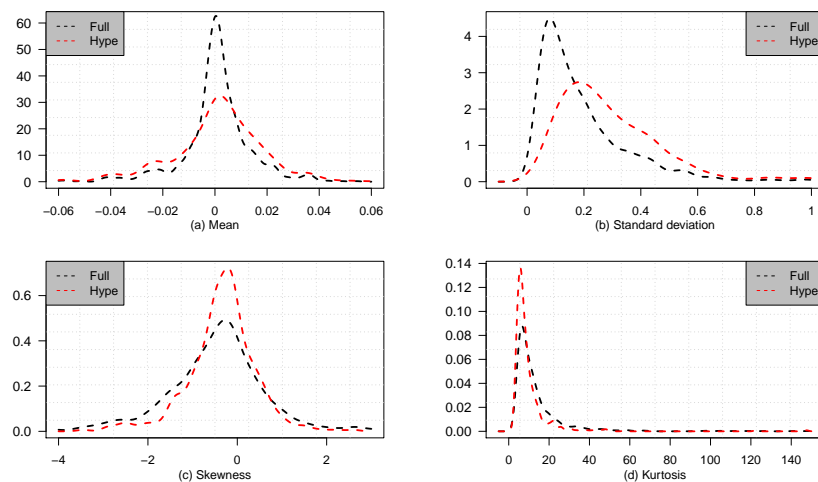
Figure 1. Gaussian kernel densities estimated using 5 min returns over the daily mean (a); standard deviation (b); skewness (c) and excess of kurtosis (d) coefficients for Bitstamp (panel **A**) and Coinbase (panel **B**). Results are reported for the full sample (black) and for the Hype period (red).

Table 4. Summary statistics of the 5 min Bitcoin log returns. Results are reported for the two exchanges Bitstamp and Coinbase over the full period "Full" and conditional on the Hype period.

	Bitstamp		Coinbase	
	Full	Hype	Full	Hype
Maximum	61.09	7.41	10.62	10.62
Minimum	−36.89	−15.54	−21.02	−21.02
Mean	0.00	0.00	0.00	0.00
Median	0.00	0.00	0.00	0.00
Std. Dev.	0.31	0.33	0.20	0.31
Skewness	−0.28	−0.38	−0.58	−0.40
Excess of kurtosis	8.84	8.28	14.36	10.22

3.1. Are Bitcoin Returns Predictable?

Catania et al. (2019) investigate the predictability of Cryptocurrencies returns—and in particular Bitcoin—at one–day horizon.[3] They find evidence of predictability for Bitcoin returns at the one–day frequency when averaging over a large number of Dynamic Linear Models resorting to the Dynamic Model Averaging technique. In this section, we only focus on the plain autoregressive model of order one defined as:

$$\tilde{r}_{j(N),t} = \mu_N + \phi_N \tilde{r}_{j(N)-1,t} + \sigma \varepsilon_{j(N),t}, \quad \varepsilon_{j(N),t} \overset{iid}{\sim} (0,1),$$

where ϕ_N is the first order autoregressive coefficients for frequency N. Figure 2 plots the estimated coefficient ϕ_N for Bitstamp and Coinbase according to different values of N starting from $N = 300$ (five minutes) to $N = 2,592,000$ (30 days). Results for the Hype period are also reported. We find that ϕ_N is negative and statistical different from zero when returns are aggregated up to 6 h. The autoregressive coefficient follows an upward trend and a peculiar curve around the 12 h aggregation frequency. After this point, the estimated coefficient decreases again and start being quite noisy around 0. We find that this behaviour is consistent across exchanges and also holds during the Hype sub-sample.

We conclude that there is no strict evidence indicating that Bitcoin returns can be predicted using a first order autoregressive model when looking at horizons longer than a day. However, by looking at intraday horizons, and especially within the first 6 hours, it seems like there is some predictability, even though the statistical significance is limited.[4]

Gencay et al. (2001) note that the first order autocorrelation of high frequency financial assets is time–varying resulting in different patterns of predictability over time. We follow their approach and investigate the stability of the estimated ϕ_N over time for different N. We expect that due to the increase in the number of transactions the Bitcoin market has become more efficient over time, resulting in insignificant predictability based on prior observations. Thus, we estimate ϕ_N for $N = 900$ (15 min), $N = 1800$ (30 min), and $N = 3600$ (one hour) using only observations available in the previous month of data and update its value according to a rolling window of fixed length. Figure 3 displays the estimated ϕ_N coefficients for Bitstamp and Coinbase. We find that during the begin of the sample Bitstamp shows a significant predictability pattern for both the 30 min and one hour intervals. Though, from the beginning of 2015, this predictability seems to have shrinkage down and lead into the 95% confidence range indicating insignificant predictability. This pattern indicates that Bitcoin traded at Bitstamp has become more efficient especially during the Hype period. Differently, in the Coinbase exchange we do not find the same predictability pattern. A possible explanation could be the more

---

[3]　See also Balcilar et al. (2017) who examine the causal relation between Bitcoin return/volatility and traded volumes.

[4]　However, we acknowledge that at the time of writing there is a lag of 10.83 min between the placement and execution of a trade on Bitstamp. Differently, on Coinbase trade execution is immediate.

substantial amount of trades for the Coinbase exchange compared to Bitstamp, which implies that Coinbase is a more liquid and efficient market.
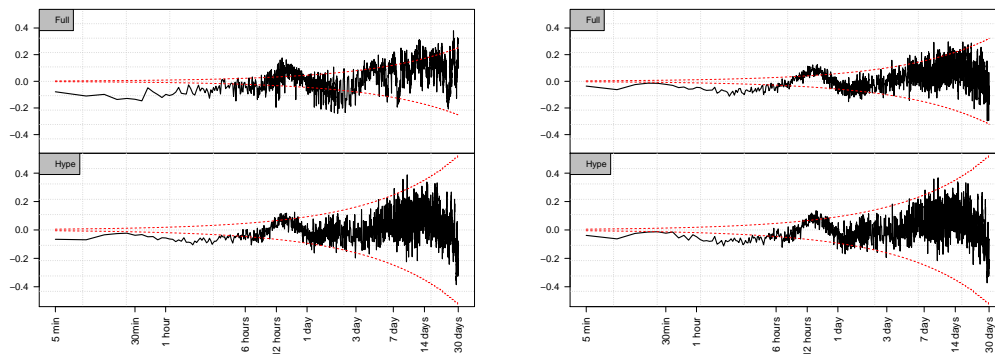
Figure 2. Linear correlation coefficient calculated as a function of the size of the time interval of returns for Bitstamp left and Coinbase right. The horizontal axis is the logarithmic of the time interval.
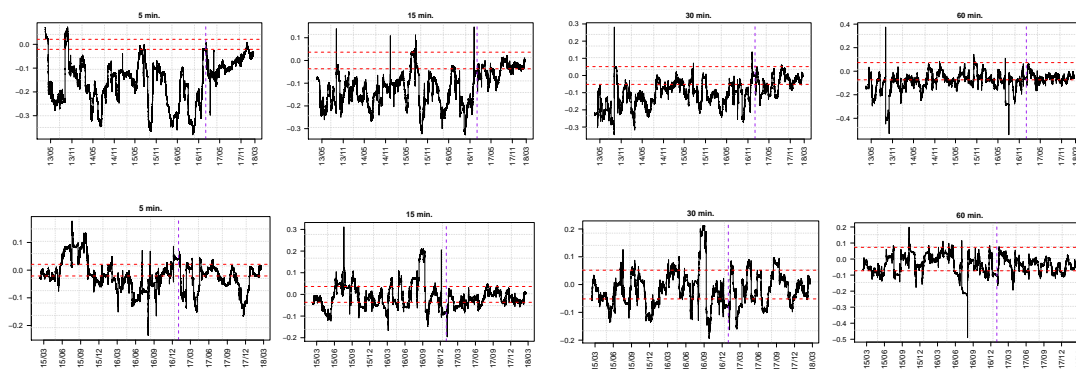

Figure 3. First order serial autocorrelation coefficient estimated using a fixed rolling windows of one month for Bitstamp (top figures) and Coinbase (bottom figures). Horizontal dashed lines indicate the 95% confidence interval. The dashed vertical line indicates the start of the Hype period at the begin of January 2017. Results are reported for Bitcoin logarithmic returns sampled at 5, 15, 30, and 60 min.

3.2. Seasonality in Bitcoin's Volatility

Similar to foreign exchange rates, also Bitcoin exhibits a large amount of seasonality in its volatility, see for example Dacorogna et al. (1993), Taylor and Xu (1997), and Breedon and Ranaldo (2013). We investigate the daily seasonality pattern by looking at the intraday realized volatility computed at 30 min over the full sample and over the Hype period, as well as the average traded volumes computed at the same frequency.

Figure 4 reports a graphical illustration of the intraday realized volatility and average volumes every 30 min. The figure shows a clear seasonality pattern for the average traded volume (vertical lines) and intraday realized volatility (red line). It is interesting to see the differences in the two exchanges peak hours, referring to the fact that Bitstamp is a European-based exchange, and Coinbase is a US-based exchange. Hence, they have spikes in different timezones associated with their working hours. Asia should be represented in both exchanges but does not seem to influence the two figures. We also note that the seasonal pattern has not changed during the Hype period.
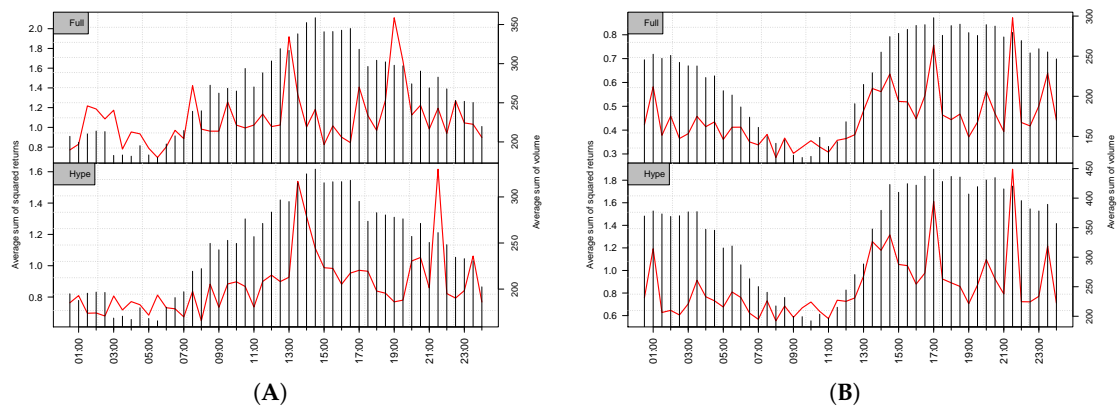
Figure 4. Intraday realized volatility (red lines, left axis) and average volumes (vertical bars, right axis) computed every 30 min for Bitcoin traded at the Bitstamp (panel **A**) and Coinbase (panel **B**) exchanges.

Besides the analysis of the intra-daily seasonality, we also investigate whether there is presence of intra-weekly seasonality. Following Dacorogna et al. (1993), we divide the week into a sequence of 2 h equally spaced observations. Figure 5 reports the weekly sequence based on average realized volatility (red line) and the average sum of volume (vertical lines). The effect of the weekends is clear from the figure. However, we note that this effect is more pronounced during the Hype period reported in the top panel of the figure. Interestingly, we also find evidence of intra–weekly seasonality for other days. Indeed, for both the exchanges we observe increasing activity from Monday to Thursday–Friday and then a decreasing curve over Saturday and Sunday. Remarkably, this effect during working days is not present for regular financial trading assets.
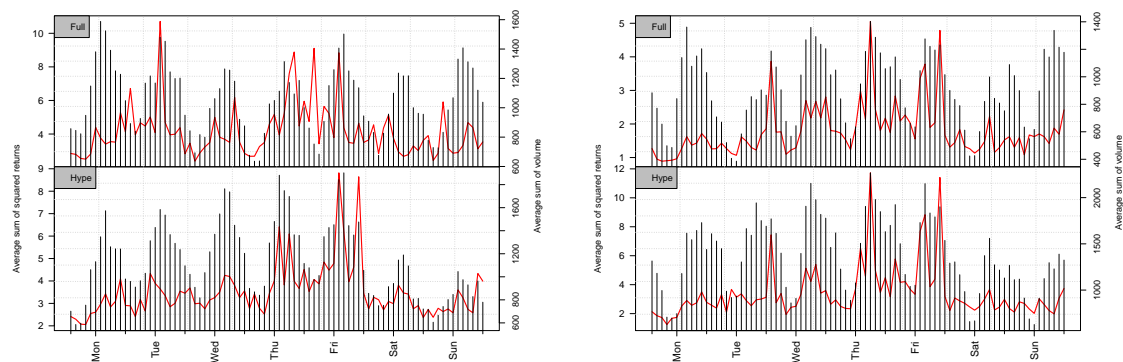
Figure 5. Weekly seasonality computed for the Bitstamp (left figures) and Coinbase (right figures) exchanges over the full sample and Hype period. Red lines report the average realized volatility (**left axis**) while the vertical bars report the average sum of volume (**right axis**).

4. Modelling and Predicting Bitcoin Realized Volatility

In this section, we report on an in sample and out of sample forecast analysis of the Bitcoin's realized volatility using several Heterogeneous Autoregressive (HAR) specifications. HAR has been originally introduced by Corsi (2008) in order to approximate the slow decay of the autocorrelation function of realized volatility. The model builds on the assumption of three different types of investors creating three different types of volatility. The investors are: (i) short-term traders with daily activity; (ii) medium investors who typically regulate their portfolio once a week; and (iii) long-term investors with horizon around a month or longer. Corsi (2008) and Corsi et al. (2012) argue that while the level of short-term volatility does not affect the long-term traders, the level of long-term volatility does affect the short-term traders, as it determines the expectation to the future size of trends and risks. Hence,

the short-term volatility is dependent on the longer horizon volatility, while the long-term volatility only consist of an $AR\,(1)$ structure, then the model can be written in a hierarchical system defined by

$$
\begin{aligned}
\tilde{\sigma}^m_{t+1m} &= c^m + \phi RV^m_t + \tilde{\omega}^m_{t+1m} \\
\tilde{\sigma}^w_{t+1w} &= c^w + \phi^w RV^w_t + \gamma^w E_t \left[\tilde{\sigma}^m_{t+m} \right] + \tilde{\omega}^w_{t+1m} \\
\hat{\sigma}^d_{t+1d} &= c^d + \phi^d RV^d_t + \gamma^d E_t \left[\tilde{\sigma}^w_{t+w} \right] + \tilde{\omega}^d_{t+1d}
\end{aligned}
\tag{2}
$$

where RV^d_t, RV^w_t and RV^m_t are the daily, weekly and monthly realized volatility and $\tilde{\omega}^d_{t+1d}$, $\tilde{\omega}^w_{t+1w}$, and $\tilde{\omega}^m_{t+1m}$ are the volatility innovations for the daily, weekly and monthly horizons, respectively. The economic interpretation of this hierarchical system is that each horizon volatility component consists of two parameters: (i) the expectation to the next period volatility; and (ii) an expectation for the longer horizon volatility, which is shown to have an impact on the future volatility. The HAR model can be written in a cascade of previous values for one day, one week and one month. By straightforward recursive substitutions we obtain a forecasting model for the realized volatility as:

$$
RV^d_{t,t+h} = c + \beta^d RV^d_t + \beta^w RV^w_t + \beta^m RV^m_t + \epsilon_{t,t+h}
\tag{3}
$$

where $h \geq 0$ is the forecast horizon and $\epsilon_{t,t+h}$ is a zero mean serially uncorrelated shocks and $RV^w_t = \frac{1}{7} \sum_{s=1}^{7} RV^d_{t-s+1}$, and $RV^m_t = \frac{1}{28} \sum_{s=1}^{28} RV^d_{t-s+1}$ are the weekly and monthly volatility, respectively.[5] This model is labelled as "HAR-RV". Andersen et al. (2007) extended the HAR-RV model to include a jump component in the cascade of lagged volatility measures. Jumps are defined as:

$$
J_{t+1} \equiv \max \left(RV_{t+1} - BV_{t+1}, 0 \right)
\tag{4}
$$

where:

$$
BV_{t+1} = \mu_1^{-2} \sum_{i=2}^{n} |r_{t+i}| \left| r_{t+(i-1)} \right|,
\tag{5}
$$

with $\mu_1 = \sqrt{2/\pi}$ is the bipower variation introduced by Barndorff-Nielsen and Shephard (2003) and Barndorff-Nielsen (2004). By including the jump component into the HAR model we obtain the "HAR-RV-J" defined as:

$$
RV^d_{t,t+h} = c + \beta^d RV^d_t + \beta^w RV^w_t + \beta^m RV^m_t + \alpha^d J^d_t + \epsilon_{t,t+h}.
\tag{6}
$$

A related specification has been further introduced by Barndorff-Nielsen et al. (2006) by including the so called "significant jumps" component. Specifically, let:

$$
TQ_{t+1} = n\mu_{\frac{4}{3}}^{-3} \sum_{i=3}^{n} |r_{t,i}|^{\frac{4}{3}} |r_{t,i-1}|^{\frac{4}{3}} |r_{t,i-2}|^{\frac{4}{3}}
\tag{7}
$$

be the realized tripower quarticity where $\mu_{\frac{4}{3}} = 2^{\frac{2}{3}} \cdot \Gamma\,(7/6) \cdot \Gamma\,(1/2)^{-1}$. The significant jump component at level $\tau \in (0,1)$ is defined as:

$$
J_{t+1,\tau} = I\left[Z_{t+1} > \Phi_{1-\tau} \right] \cdot \left[RV_{t+1} - BV_{t+1} \right]
\tag{8}
$$

where $I\,[A]$ is the indicator function equal to 1 if A is true and 0 otherwise, and:

$$
Z_{t+1} = \Delta^{-1/2} \times \frac{\left[RV_{t+1} - BV_{t+1} \right] RV_{t+1}^{-1}}{\left[\left(\mu_1^{-4} + 2\mu_1^{-2} - 5 \right) \max \left\{ 1, TQ_{t+1} BV_{t+1}^{-2} \right\} \right]^{-\frac{1}{2}}}
\tag{9}
$$

---

5   Please note that Bitcoin is traded 7 days a week.

is the feasible test statistics arising from the asymptotic distribution of the difference between the realized volatility and the bipower variation, see Barndorff-Nielsen et al. (2006) for more details. Finally, the new HAR model with continuous jumps, HAR-RV-CJ, is defined as:

$$RV_{t,t+h}^d = c + \beta^{cd} C_t^d + \beta^{cw} C_t^w + \beta^{cm} C_t^m + \alpha^{cd} J_{t,\tau}^d + \alpha^{cw} J_{t,\tau}^w + \alpha^{cm} J_{t,\tau}^m + \epsilon_{t,t+h} \tag{10}$$

where:

$$C_{t+1} = I\left[Z_{t+1} \leq \Phi_{1-\tau}\right] RV_{t+1} + I\left[Z_{t+1} > \Phi_{1-\tau}\right] BV_{t+1} \tag{11}$$

selects RV_{t+1} if $Z_{t+1} \leq \Phi_{1-\tau}$ and BV_{t+1} if $Z_{t+1} > \Phi_{1-\tau}$. We perform a sensitivity analysis similar to that reported in Andersen et al. (2007) and set $\tau = 0.01$.

4.1. Including a Leverage Component

A well known stylized fact of equity financial returns is the so called leverage effect, see Black (1976), Nelson (1991) and Zakoian (1994), among others. The leverage effect relates to the different reaction of the volatility of a firm to past positive and negative news. Its original formulation relates to the reaction of the volatility to changes in the debt to equity ratio of a traded company. Specifically, when a bad news arrives, the value of the firm decreases while its debt remains unchanged. This leads to an increase of the debt to equity ratio corresponding to an increase of the riskiness of the firm which translates in more volatility. Of course, the original interpretation of the leverage effect cannot be applied to Bitcoin since it does not have any capital structure. However, previous empirical works have found evidence of leverage effect for Bitcoin, see Catania and Grassi (2017), Katsiampa (2017), Bariviera (2017), and Ardia et al. (2018). We follow Corsi et al. (2012) and introduce a leverage component in the HAR specification by defining:

$$r_t^{-d} = \min\left[\sum_{i=1}^{288} r_{t,i}, 0\right] \tag{12}$$

which indicates the minimum return over the trading day. The variable r_t^{-d} along with its weekly r_t^{-w} and monthly r_t^{-m} averages are included linearly in the HAR-RV, HAR-RV-J, and HAR-RV-CJ specifications. For example, the HAR-RV specification with leverage, HAR-RV-L, is defined as:

$$RV_{t,t+h}^d = c + \beta^d RV_t^d + \beta^w RV_t^w + \beta^m RV_t^m + \gamma^d r_t^{-d} + \gamma^w r_t^{-w} + \gamma^m r_t^{-m} + \epsilon_{t,t+h} \tag{13}$$

4.2. In Sample Results

We consider the realized variance of Bitcoin from 17 March 2013 for Bitstamp and from 2 February 2015 for Coinbase up to 18 March 2018. Similar to previous results, we also consider the Hype period from 1 January 2017 to 18 March 2018. Results are also reported for the realized standard deviation, $RSD = \sqrt{RV}$ and the logarithmic realized variance, $LRV = \log(RV)$. Figure 6 displays: (i) the time series of log realized variance; (ii) the feasible test statistics; and (iii) the significant logarithmic jump series, $\log(J_{t,\tau} + 1)$ over the full sample for Bitstamp and Coinbase. We find that the logarithmic realized variance for the Coinbase exchange displays an increasing pattern, with the highest values in the end, and especially around December 2017 where the underline value increased significantly. Interestingly, we find that realized volatility is lower during the bubble period of 2017 compared to the bear market period of 2018.

Panel (*b*) reports the test statistics from Equation (9) for $\tau = 0.01$. The red horizontal line indicates $\Phi_{1-0.01} = 2.32$, i.e., the threshold after which jumps are classified as significant. Interestingly, we find a very large proportion of jumps for Bitcoin compared to the proportion usually found in other asset classes, see e.g., Andersen et al. (2007). Indeed, the proportion of jumps ranges from 27% to 92% depending on different choices of τ. When $\tau = 0.01$, the proportion of jumps over the full period is around 79% for Bitstamp and 85% for Coinbase. If we focus on the Hype period the proportion of

jumps is halved for both exchanges. This results further indicates the growing trade intensity and the increased stability of the market over time.
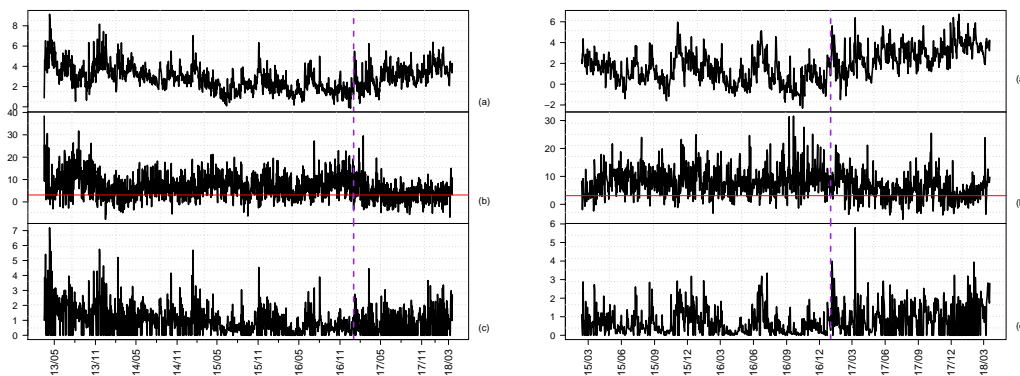
Figure 6. Plot of logarithmic realized variance $\log(RV_t)$ (**a**), Z_t (**b**) and logarithmic significant jumps $\log(J_t + 1)$ (**c**) over time. Purple vertical dashed lines indicate the start of the Hype period. The horizontal red line indicates the $1 - \tau$ quantile of a standard Gaussian distribution for $\tau = 0.01$. Figures on the left panel are for Bitstamp, figures on the right panel for Coinbase.

Table 5 reports the summary statistics for the realized variance and its transformations. We find that both the median and the standard deviation of the realized variance and jump component are higher during the Hype period. We also find that similar to Andersen et al. (2001a) and Andersen et al. (2001b), we are not able to reject the null hypothesis of normality for the logarithmic realized variance according to the Jarque-Bera test statistics.

Table 5. Summary statistic for the realized variance, realized standard deviation and logarithmic realized variance. Panel (A) reports results for the full sample while panel (B) for the Hype period. The rwo J.test reports the Jarque-Bera test statistics for the null hypothesis of Gaussianity.

Panel (A)—Full sample

	Bitstamp						Coinbase					
	RV_t	$RV_t^{\frac{1}{2}}$	$\log(RV_t)$	J_t	$J_t^{\frac{1}{2}}$	$\log(J_t+1)$	RV_t	$RV_t^{\frac{1}{2}}$	$\log(RV_t)$	J_t	$J_t^{\frac{1}{2}}$	$\log(J_t+1)$
Maximum	9374.99	96.82	9.15	1320.43	36.34	7.19	835.73	28.91	6.73	326.73	18.08	5.79
Minimum	0.87	0.93	−0.14	0.12	0.35	0.12	0.01	0.31	−2.31	0.02	0.15	0.02
Mean	54.94	5.23	2.86	7.35	1.85	1.29	21.47	3.41	1.82	2.69	1.25	0.88
Median	15.63	3.95	2.75	2.01	1.42	1.10	5.81	2.41	1.76	1.02	1.00	0.70
Std. Dev.	311.15	5.26	1.23	45.42	1.98	0.87	53.93	3.14	1.57	11.48	1.06	0.73
Skewness	21.19	7.39	0.69	21.16	7.96	1.82	7.51	2.72	0.19	24.81	5.51	1.34
Kurtosis	541.82	94.59	4.11	552.17	102.12	8.60	81.77	14.48	2.55	694.19	74.09	5.77
J.test	$22e^{+6}$	$65e^{+4}$	239				$30e^{+4}$	7612	16.01			

Panel (B)—Hype period

	Bitstamp						Coinbase					
	RV_t	$RV_t^{\frac{1}{2}}$	$\log(RV_t)$	J_t	$J_t^{\frac{1}{2}}$	$\log(J_t+1)$	RV_t	$RV_t^{\frac{1}{2}}$	$\log(RV_t)$	J_t	$J_t^{\frac{1}{2}}$	$\log(J_t+1)$
Maximum	588.38	24.26	6.38	37.84	6.15	3.66	835.73	28.91	6.73	326.73	18.08	5.79
Minimum	1.40	1.19	0.34	0.26	0.51	0.23	0.23	0.48	−1.48	0.07	0.26	0.06
Mean	42.23	5.64	3.18	4.62	1.91	1.44	41.04	5.21	2.89	5.16	1.78	1.30
Median	24.56	4.96	3.20	2.90	1.70	1.36	18.70	4.32	2.93	2.19	1.48	1.16
Std. Dev.	58.99	3.23	1.06	5.37	0.99	0.72	76.98	3.73	1.29	21.39	1.42	0.76
Skewness	4.44	1.80	0.06	3.10	1.32	0.53	5.49	2.30	−0.06	14.05	6.77	1.41
Kurtosis	31.60	8.08	2.86	16.08	5.49	2.90	43.11	10.97	3.16	210.54	73.42	7.91
J. test	16,521	714	0.62				31,853	1562	0.67			

Model Estimation

We now estimate by OLS the HAR-RV, HAR-RV-J, and HAR-RV-CJ models to the realized variance, realize standard deviation and logarithmic realized variance over the full sample for the two exchanges.

Specifications that include the leverage component are also estimated and indicated with the additional label "-L". Estimation results are reported in Table 6. Estimated coefficients are in line with those usually found in the literature for other asset classes. Interestingly, we find that specifications that include the leverage component outperform their counterpart without leverage. Regarding the estimated leverage coefficients, we see that these are negative and statistically significant at standard confidence levels. This finding is somehow in contrast with previous results by results by Catania and Grassi (2017) and Ardia et al. (2018) who document an "inverted" leverage effect for Bitcoin. To further investigate this aspect, in Figure 7 we report the empirical autocorrelation at different lags between realized variance and the leverage component, i.e., $cor(RV_t, r_{t-h}^{-d})$ for $h = 1, \ldots, 50$. The plot is reported for the two exchanges for the full sample as well as for the Hype period. Results indicate that correlations are negative and statistically different from zero up to $h = 10$ when computed over the full sample. However, when we focus on the Hype period, evidence of correlation between RV_t and r_{t-h}^{-d} is less strong. This result suggests that the leverage effect has changed over time for Bitcoin and somehow confirms the findings of Ardia et al. (2018).

(A)   (B)   (C)   (D)
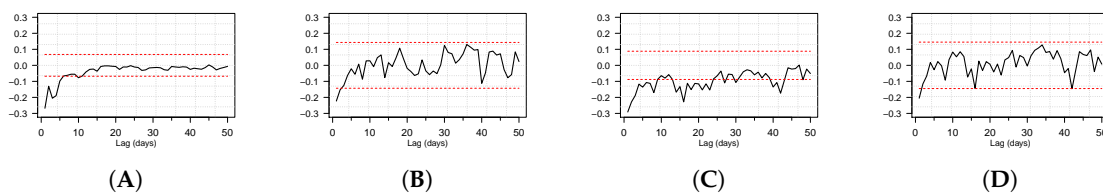
Figure 7. Empirical cross correlation at different lags between realized variance and the leverage component, $cor(RV_t, r_{t-h}^{-d})$ for $h = 1, \ldots, 50$. Panels (**A**) and (**B**) report results for Bitstamp over the full sample and the Hype period, respectively. Panels (**C**) and (**D**) report results for Coinbase over the full sample and the Hype period, respectively. Horizontal red dashed lines indicate 95% confidence bounds.

4.3. Out of Sample Results

We now conduct an out of sample analysis studying the predictability of Bitcoin realized variance at different horizons. Predictions are made by the models previously introduced at horizons $h = 1$ (one day), $h = 7$ (one week), and $h = 28$ (one month) using the direct method of forecast, see Marcellino et al. (2006). We start making prediction from 21 April 2014 for Bitstamp and 17 March 2016 for Coinbase and than update model parameters each time a new observation becomes available during the whole forecast periods using a fixed rolling window. The length of the out of sample is $F = 1424$ and $F = 731$ for Bitstamp and Coinbase, respectively. Results are compared with the Random Walk (RW) specification defined by:

$$RV_{t+h} = RV_t + \epsilon_{t+h}. \tag{14}$$

Let $\widehat{RV}_{t+h}$ be the prediction made at time t for time $t + h$. Comparison among different specification is performed according to the mean absolute forecast error (MAFE) and root mean square forecast error (RMSFE). MAFE at horizon h is defined as:

$$MAFE_h = \frac{1}{F} \sum_{f=1}^{F} \left| RV_{T+f} - \widehat{RV}_{T+f+h} \right|, \tag{15}$$

while RMSFE as:

$$RMSFE_h = \sqrt{\frac{1}{F} \sum_{f=1}^{F} \left(RV_{T+f} - \widehat{RV}_{T+f+h} \right)^2} \tag{16}$$

where T is the length of the in sample period. Models with lower MAFE and RMSFE are preferred. Table 7 reports the results computed over the full sample. Results for the Hype period are similar and are available upon request to the second author.

Table 6. OLS in sample estimated coefficients for the five different HAR models for Bitstamp (panel A) and Coinbase (panel B). HAC standard errors are reported in parentheses based on the Newey-West correction. The apexes *a–d* report the significance level, $a = 0.1\%$, $b = 1\%$, $c = 5\%$ and $d = 10\%$.

Panel (A)—Bitstamp

	HAR-RV			HAR-RV-J			HAR-RV-L			HAR-RV-CJ			HAR-RV-CJ-L		
	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$
c	17.15^a (2.20)	1.07^a (0.14)	0.24^a 0.05	17.16^a (2.22)	1.09^a (0.15)	0.23^a 0.05	13.96^c (6.12)	1.24^a (0.13)	0.18^b 0.06	16.92^a (1.92)	1.13^a (0.15)	0.40^a (0.05)	15.01^c (6.93)	1.30^a (0.14)	0.34^a (0.06)
β^d	0.24^a (0.07)	0.47^a (0.05)	0.57^a 0.03	0.24^a (0.06)	0.48^a (0.05)	0.59^a 0.03	0.15 (0.11)	0.37^a (0.09)	0.51^a 0.03	0.28^a (0.05)	0.49^a (0.05)	0.53^a (0.03)	0.17^c (0.08)	0.38^a (0.08)	0.47^a (0.03)
β^w	0.16^c (0.07)	0.12^d (0.07)	0.18^a 0.04	0.16^c (0.07)	0.12^d (0.07)	0.18^a 0.04	0.14 (0.09)	0.08 (0.09)	0.20^a 0.05	0.09 (0.11)	0.09 (0.08)	0.20^a (0.05)	0.07 (0.11)	0.04 (0.11)	0.23^a (0.05)
β^m	0.13^b (0.04)	0.16^a (0.04)	0.14^a 0.04	0.13^b (0.04)	0.16^a (0.04)	0.14^a 0.03	0.16^b (0.06)	0.28^a (0.08)	0.21^a 0.05	0.24 (0.15)	0.19^b (0.07)	0.14^b (0.04)	0.31 (0.19)	0.37^b (0.11)	0.21^a (0.05)
α^d			−0.02 (0.56)	−0.07 (0.12)	−0.04 0.03			−0.15 (0.86)	−0.01 (0.14)	0.04 (0.03)	0.01 (0.69)	0.01 (0.12)	0.04 (0.02)		
α^w							0.74 (1.17)	0.13 (0.21)	−0.04 (0.05)	0.75 1.02	0.13 (0.20)	−0.06 (0.05)			
α^m							−0.80 (1.11)	−0.10 (0.18)	0.00 (0.06)	−0.95 1.25	−0.19 (0.19)	0.03 (0.06)			
γ^d				-7.41^d (3.91)	-0.19^c (0.08)		-0.03^a 0.01			-7.28^d 3.75	-0.18^c (0.08)	-0.03^a (0.01)			
γ^w				−3.50 (4.70)	−0.16 (0.14)		0.01 0.02			−4.18 4.55	−0.18 (0.14)	0.0 (0.02)			
γ^m				6.54 (4.58)	0.49 (0.22)		0.08^c 0.03			8.57 5.61	0.55^c (0.24)	0.10^b (0.03)			
R^2	23%	50%	67%	23%	50%	67%	25%	51%	68%	23%	50%	67%	25%	51%	68%
adj. R^2	23%	49%	67%	23%	49%	67%	25%	51%	68%	23%	49%	67%	25%	51%	67%

Panel (B)—Coinbase

	HAR-RV			HAR-RV-J			HAR-RV-L			HAR-RV-CJ			HAR-RV-CJ-L		
	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$
c	4.44^b (1.70)	0.40^a (0.12)	0.02 (0.06)	5.01^b (1.92)	0.44^c (0.17)	0.05 (0.06)	8.06^c (3.67)	0.44^a (0.12)	0.01 (0.06)	6.67^a (1.81)	0.65^a (0.17)	0.30^a (0.06)	8.69^c (3.69)	0.68^a (0.18)	0.29^a (0.07)
β^d	0.33^a (0.10)	0.51^a (0.06)	0.56^a (0.03)	0.37^b (0.12)	0.52^a (0.06)	0.59^a (0.04)	0.12^c (0.12)	0.42^a (0.07)	0.52^a (0.04)	0.37^b (0.12)	0.51^a (0.06)	0.50^a (0.36)	0.31^c (0.14)	0.43^a (0.08)	0.45^a (0.04)
β^w	0.08 (0.07)	0.13^b (0.05)	0.24^a (0.04)	0.08 (0.07)	0.15^b (0.05)	0.25^a (0.04)	−0.05 (0.12)	0.11^d (0.06)	0.27^a (0.05)	0.04 (0.09)	0.14^d (0.07)	0.31^a (0.06)	−0.07 (0.12)	0.11 (0.07)	0.33^a (0.06)
β^m	0.40^c (0.18)	0.19^b (0.06)	0.11^b (0.04)	0.38^c (0.19)	0.19^b (0.06)	0.10^c (0.04)	0.69^d (0.36)	0.37^b (0.13)	0.15^b (0.06)	0.47^c (0.23)	0.25^b (0.08)	0.12^c (0.05)	0.68^d (0.37)	0.39^b (0.13)	0.16^c (0.06)
α^d			−0.46 (0.32)	−0.10 (0.18)	−0.08 (0.05)			−0.09 (0.23)	−0.12 (0.13)	0.09^d (0.05)	−0.18 (0.19)	−0.08 (0.13)	0.08^d (0.04)		
α^w							0.42 (0.33)	0.02 (0.1)	-0.16^d (0.09)	0.17 (0.34)	0.00 (0.16)	-0.16^d (0.09)			
α^m							−1.20 (0.84)	-0.27^d (0.16)	−0.05 (0.08)	−0.51 (0.53)	−0.18 (0.14)	−0.04 (0.08)			
γ^d				-3.70^b (1.18)	-0.18^a (0.05)		-0.05^a (0.01)			-3.70^b (1.18)	-0.18^a (0.05)	-0.05^a (0.01)			
γ^w				−3.72 (4.14)	−0.07 (0.16)		0.04 (0.04)			−3.94 (4.01)	−0.07 (0.16)	0.04 (0.04)			
γ^m				12.82 (8.20)	0.56^c (0.28)		0.08 (0.06)			11.01 (7.80)	0.51^d (0.27)	0.08 (0.06)			
R^2	23%	51%	68%	23%	51%	68%	26%	52%	68%	24%	51%	68%	27%	53%	68%
adj. R^2	23%	50%	68%	23%	51%	68%	26%	52%	68%	24%	51%	68%	26%	53%	68%

Table 7. This table reports the out of sample forecast results for different HAR models. Results for Bitstamp are reported in panel (A) for Coinbase in panel (B) for the three forecast horizons $h = 1$, $h = 7$, and $h = 28$. Diebold Mariano test statistics with respect to the RW model (DM1) and HAR-RV model (DM2) are reported. p-values based on the asymptotic Gaussian distribution with HAC standard errors are reported in parenthesis.

Panel (A)—Bitstamp

Daily $h = 1$

Model	RW			HAR-RV			HAR-RV-J			HAR-RV-L			HAR-RV-CJ			HAR-RV-CJ-L		
	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$
R^2	24.57%	50.18%	63.51	21.21%	51.00%	65.53%	20.79%	50.84%	65.59%	20.70%	50.48%	65.91%	20.40%	50.97%	65.53%	20.30%	50.33%	65.94%
MAFE	16.36	1.20	0.51	24.16	1.24	0.47	24.12	1.24	0.47	25.03	1.27	0.47	23.93	1.25	0.48	25.34	1.27	0.47
RMSFE	52.19	2.12	0.68	47.54	1.97	0.63	47.57	1.97	0.63	48.15	1.98	0.62	47.44	1.97	0.63	48.82	1.98	0.63
DM_1				0.93 (35.13)	1.79 (7.44)	5.59 (0.00)	0.92 (35.63)	1.79 (7.41)	5.65 (0.00)	0.87 (38.37)	1.78 (7.46)	5.87 (0.00)	0.96 (33.67)	1.85 (6.38)	5.28 (0.00)	0.82 (40.98)	1.88 (6.08)	5.65 (0.00)
DM_2							−1.27 (20.27)	−0.36 (71.53)	1.19 (23.38)	−0.31 (76.00)	−0.33 (74.51)	1.76 (7.91)	0.36 (72.11)	0.59 (55.27)	−1.00 (31.90)	−0.59 (55.67)	−0.19 (84.87)	1.11 (26.60)

Weekly $h = 7$

Model	RW			HAR-RV			HAR-RV-J			HAR-RV-L			HAR-RV-CJ			HAR-RV-CJ-L		
	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$
R^2	15.48%	33.56%	47.95%	4.70%	29.47%	46.14%	4.30%	29.71%	46.32%	0.89%	23.37%	46.90%	5.12%	30.05%	46.42%	0.26%	21.74%	46.56%
MAFE	17.86	1.45	0.59	27.04	1.62	0.59	27.14	1.62	0.59	29.52	1.62	0.56	26.65	1.62	0.59	30.36	1.64	0.57
RMSFE	36.05	2.21	0.76	35.15	2.08	0.72	35.33	2.08	0.72	38.34	2.15	0.71	34.91	2.08	0.72	39.66	2.17	072
DM_1				0.69 (48.99)	2.66 (0.78)	3.97 (0.01)	0.56 (57.57)	2.65 (0.82)	4.03 (0.01)	−1.71 (8.67)	1.13 (25.68)	4.76 (0.00)	0.88 (37.81)	2.73 (0.64)	3.87 (0.01)	−2.76 (11.50)	0.62 (53.72)	4.37 (0.00)
DM_2							−1.38 (16.71)	0.03 (97.80)	1.33 (18.52)	−8.25 (0.00)	−2.58 (1.01)	2.40 (1.66)	1.69 (9.07)	1.00 (31.84)	−0.24 (81.33)	−8.62 (0.00)	−2.95 (0.32)	1.63 (10.27)

Monthly $h = 28$

Model	RW			HAR-RV			HAR-RV-J			HAR-RV-L			HAR-RV-CJ			HAR-RV-CJ-L		
	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$	RV_t	$RV_t^{\frac{1}{2}}$	$Log(RV_t)$
R^2	19.46%	24.87%	30.15%	0.37%	14.29%	27.57%	0.37%	14.18%	27.62%	1.41%	8.20%	31.14%	1.41%	20.04%	32.86%	0.54%	9.99%	33.79%
MAFE	16.22	1.52	0.65	28.00	1.88	0.66	28.01	1.88	0.66	31.07	1.83	0.61	27.32	1.83	0.64	31.55	1.83	0.60
RMSFE	24.82	2.05	0.82	31.79	2.13	0.77	31.80	2.13	0.77	36.82	2.17	0.74	31.04	2.08	0.76	38.52	2.23	0.73
DM_1				−10.06 (0.00)	−1.84 (6.54)	3.83 (0.01)	−10.08 (0.00)	−1.87 (6.12)	3.82 (0.01)	−13.54 (0.00)	−2.38 (1.75)	7.49 (0.00)	−8.88 (0.00)	−0.70 (48.18)	5.21 (0.00)	−13.85 (0.00)	−3.48 (0.05)	8.22 (0.00)
DM_2							−2.27 (2.35)	−2.34 (1.93)	−0.71 (47.56)	−12.98 (0.00)	−1.58 (11.33)	5.13 (0.00)	8.79 (0.00)	6.53 (0.00)	5.15 (0.00)	−11.76 (0.00)	−2.94 (0.33)	6.00 (0.00)

Table 7. *Cont.*

Panel (*B*)—Coinbase

Daily $h = 1$

Model	RW RV_t	RW $RV_t^{\frac{1}{2}}$	RW $Log(RV_t)$	HAR-RV RV_t	HAR-RV $RV_t^{\frac{1}{2}}$	HAR-RV $Log(RV_t)$	HAR-RV-J RV_t	HAR-RV-J $RV_t^{\frac{1}{2}}$	HAR-RV-J $Log(RV_t)$	HAR-RV-L RV_t	HAR-RV-L $RV_t^{\frac{1}{2}}$	HAR-RV-L $Log(RV_t)$	HAR-RV-CJ RV_t	HAR-RV-CJ $RV_t^{\frac{1}{2}}$	HAR-RV-CJ $Log(RV_t)$	HAR-RV-CJ-L RV_t	HAR-RV-CJ-L $RV_t^{\frac{1}{2}}$	HAR-RV-CJ-L $Log(RV_t)$
R^2	15.90%	44.46%	67.50%	17.33%	47.51%	69.58%	17.23%	47.10%	69.52%	19.28%	48.70%	69.82%	17.52%	47.92%	69.69%	19.43%	48.83%	69.85%
MAFE	23.24	1.61	0.77	20.15	1.42	0.72	20.79	1.45	0.72	19.97	1.40	0.72	21.11	1.45	0.72	20.74	1.42	0.72
RMSFE	69.21	2.87	1.02	58.25	2.57	0.95	59.34	2.58	0.95	57.48	2.54	0.95	59.37	2.56	0.95	58.36	2.53	0.95
DM_1				2.01 (4.51)	2.99 (0.29)	3.92 (0.01)	2.04 (4.16)	2.98 (0.30)	3.90 (0.01)	2.15 (3.22)	3.34 (0.09)	3.93 (0.01)	2.04 (4.20)	2.97 (0.31)	4.20 (0.00)	2.18 (2.98)	3.26 (0.12)	4.18 (0.00)
DM_2							−1.57 (11.70)	−1.22 (22.26)	0.23 (81.61)	1.34 (18.19)	1.58 (11.57)	0.13 (89.53)	−1.57 (11.68)	0.29 (77.49)	1.32 (18.68)	−0.11 (91.59)	1.07 (28.36)	0.91 (36.36)

Weekly $h = 7$

Model	RW RV_t	RW $RV_t^{\frac{1}{2}}$	RW $Log(RV_t)$	HAR-RV RV_t	HAR-RV $RV_t^{\frac{1}{2}}$	HAR-RV $Log(RV_t)$	HAR-RV-J RV_t	HAR-RV-J $RV_t^{\frac{1}{2}}$	HAR-RV-J $Log(RV_t)$	HAR-RV-L RV_t	HAR-RV-L $RV_t^{\frac{1}{2}}$	HAR-RV-L $Log(RV_t)$	HAR-RV-CJ RV_t	HAR-RV-CJ $RV_t^{\frac{1}{2}}$	HAR-RV-CJ $Log(RV_t)$	HAR-RV-CJ-L RV_t	HAR-RV-CJ-L $RV_t^{\frac{1}{2}}$	HAR-RV-CJ-L $Log(RV_t)$
R^2	15.88	38.10%	56.62%	19.01%	37.90%	57.01%	17.87%	37.90%	57.43%	20.18%	41.79%	59.80%	18.79%	38.88%	56.77%	18.09%	41.08%	58.73%
MAFE	23.85	1.72	0.90	20.02	1.73	0.84	20.29	1.73	0.83	20.33	1.67	0.82	20.82	1.72	0.84	21.67	1.69	0.82
RMSFE	42.79	2.47	1.10	36.28	2.48	1.04	36.94	2.48	1.04	36.28	2.39	1.02	37.22	2.46	1.04	38.06	2.40	1.02
DM_1				3.87 (001)	3.82 (0.01)	2.83 (0.48)	3.42 (0.07)	3.85 (0.01)	3.25 (0.12)	3.97 (0.01)	4.90 (0.00)	3.50 (1.61)	3.39 (0.07)	4.27 (0.00)	3.25 (0.12)	2.84 (0.47)	4.93 (0.00)	3.66 (0.03)
DM_2							−1.52 (12.90)	−0.33 (74.46)	1.93 (5.38)	0.00 (99.78)	1.97 (4.88)	1.61 (10.76)	−1.91 (5.69)	0.79 (42.80)	1.02 (31.04)	−1.96 (5.02)	1.57 (11.67)	1.67 (9.49)

Monthly $h = 28$

Model	RW RV_t	RW $RV_t^{\frac{1}{2}}$	RW $Log(RV_t)$	HAR-RV RV_t	HAR-RV $RV_t^{\frac{1}{2}}$	HAR-RV $Log(RV_t)$	HAR-RV-J RV_t	HAR-RV-J $RV_t^{\frac{1}{2}}$	HAR-RV-J $Log(RV_t)$	HAR-RV-L RV_t	HAR-RV-L $RV_t^{\frac{1}{2}}$	HAR-RV-L $Log(RV_t)$	HAR-RV-CJ RV_t	HAR-RV-CJ $RV_t^{\frac{1}{2}}$	HAR-RV-CJ $Log(RV_t)$	HAR-RV-CJ-L RV_t	HAR-RV-CJ-L $RV_t^{\frac{1}{2}}$	HAR-RV-CJ-L $Log(RV_t)$
R^2	32.15%	38.57	42.45	31.06%	41.34%	45.15%	30.22%	41.10%	45.06%	36.82%	46.32%	45.13%	33.57%	40.18%	37.76%	33.18%	42.74%	35.17%
MAFE	18.40	1.81	0.93	16.95	1.74	0.93	16.98	1.73	0.92	16.85	1.69	0.95	17.09	1.71	0.95	17.38	1.72	0.98
RMSFE	27.10	2.30	1.14	26.69	2.26	1.10	26.68	2.24	1.09	25.22	2.15	1.12	25.65	2.18	1.12	25.51	2.13	1.16
DM_1				0.51 (61.08)	0.71 (47.67)	1.64 (10.23)	0.52 (60.24)	1.00 (31.76)	2.04 (4.22)	2.59 (0.97)	2.48 (1.32)	0.86 (39.03)	2.07 (3.85)	2.66 (0.81)	0.95 (34.28)	2.41 (1.62)	2.97 (0.31)	−0.55 (58.02)
DM_2							0.06 (95.32)	2.40 (1.68)	3.87 (0.01)	3.91 (0.01)	2.61 (0.91)	−0.73 (46.87)	3.98 (0.01)	4.69 (0.00)	−2.13 (3.39)	2.07 (3.92)	2.84 (0.46)	−2.55 (1.09)

Along with the MAFE and RMSFE measures, the table also reports the R^2 of the Mincer–Zarnowitz regression defined by:

$$RV_{t+h} = c + \beta \widehat{RV}_{t+h} + \epsilon_{t+h}, \tag{17}$$

as well as the Diebold and Mariano (1994) test statistics of each model with respect to the benchmark RW (DM1) and with respect to the plain HAR-RV model (DM2). Results indicate that predictability is higher for lower forecast horizons. Indeed, looking at the R^2 we find that when $h = 1$, up to 65% of the log realized variance variability can be predicted with the HAR-RV-CJ model. However, when $h = 28$ the R^2 decreases to only 33%. Overall, the inclusion of jumps does not always translate in better predictions. In this respect, results are a bit mixed. Differently, models that include the leverage component seem to generally perform better than the standard HAR-RV model. Looking at the Diebold Mariano test statistic with respect to the benchmark model (DM1), we find strong evidence of predictability of all specifications. Differently, when we focus on predictability with respect to the plain HAR-RV model, results are mixed and do not show a clear pattern. Comparing results between the two exchanges indicates that realized variance is easier to predict in the Coinbase exchange.

To conclude our analysis we study the stability of prediction gains with respect to the RW benchmark over time. To do so, we compute the cumulative absolute error of a forecast model over the cumulative absolute error of the benchmark model. Specifically, the ratio of cumulative absolute errors RCAE at time f is defined as:

$$RCAE_f = \frac{\sum_{s=1}^{f} |e_{j,s}|}{\sum_{s=1}^{f} |e_{i,s}|}, \tag{18}$$

where $e_{j,s}$ is the forecast error of generic model j at time s and $e_{i,s}$ is the forecast error of the benchmark specification. Results are reported for $i = \text{RW}$ and $i = \text{HAR-RV}$. Values of $RCAE_f$ below one indicate outperformance with respect to the benchmark and viceversa. Figure 8 displays the $RCAE_f$ for the log realized variance for different forecast horizons and the two exchanges. In the top graph of each sub-figure the comparison is performed with respect to RW, while in the bottom graphs we use HAR-RV as the benchmark. Results are very clear and show that predictability of the realized variance is increased over time. Indeed, at the start of the sample we observe large losses of all models with respect to RW and HAR-RV probably due to uncertainty in estimated parameters. However, at the end of the forecasting period those losses seem to vanish suggesting that volatility becomes more easy to predict. Across the different specifications we observe that HAR-L and HAR-CJ-L are the top performer. This result confirms the in sample findings and indicates that the leverage component is important for volatility prediction of Bitcoin. A comparison across the two exchanges also suggests that volatility in the Coinbase exchange is easier to predict.
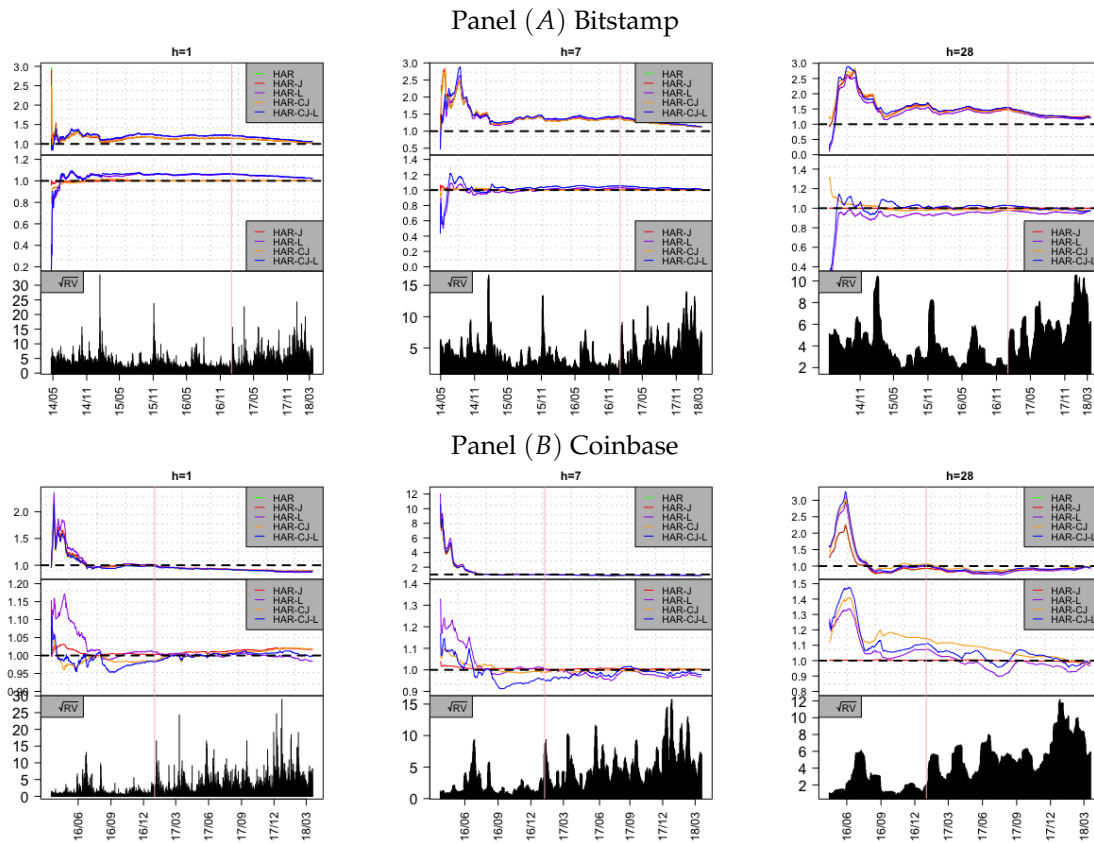
Figure 8. Relative cumulative absolute errors of several forecasting models with respect to RW (first sub-figure) and HAR-RV (second sub-figure). The third sub-figure reports the evolution of the realized standard deviation over time. The red vertical lines correspond to the start of the Hype period. Results are reported for Bitstamp in panel (**A**) and for Coinbase in panel (**B**) for the three forecast horizons $h = 1$, $h = 7$, and $h = 28$.

5. Conclusions

In this paper, we analysed Bitcoin returns sampled at high frequency and its realized variance. Raw Bitcoin transactions have been downloaded from the two exchanges Bitstamp and Coinbase. After detailing how raw data are cleaned, we started our analysis focusing on the in sample properties of Bitcoin logarithmic returns sampled at different frequencies. Results about the autocorrelation structure of Bitcoin returns as well as the intraday and intraweek seasonality of Bitcoin volatility and volumes are reported. The second part the paper focuses on predicting realized variance for Bitcoin using several forecasting models. Our results indicate that the predictability of Bitcoin realized variance is increased over time, and that predictability varies with the forecast horizon. Results extend those reported by Catania et al. (2019) and Catania et al. (2018) to the high frequency case and are consistent with previous findings reported for lower frequencies. We have also documented the presence of leverage effect for Bitcoin when realized variance is used. However, this finding is in contrast with previous results reported by Catania and Grassi (2017) and Ardia et al. (2018) where an "inverted" leverage effect is found. However, we note that differently from Catania and Grassi (2017) and Ardia et al. (2018) who use financial econometrics models to filter the conditional volatility of Bitcoin, in this study we estimate volatility using the realized variance estimator of Andersen et al. (2001a). Furthermore, our results also indicate that the leverage effect is less evident during the Hype period. Through the paper, all results have been detailed with respect to the two exchange rates as well as with a focus on the recent 2017–2018 period.

We believe that our results can be used by private investors as well as by hedge funds to improve their forecasting models for Bitcoin as well as for pricing of derivative securities. Furthermore, central banks who are considering issuing their own digital currency like Ecuador, Tunisia, and Sweden can exploits our results to improve the efficiency and reduce the volatility of the resulting market.

We have not investigated the presence of possible arbitrage opportunities across the two exchanges. Possible extensions can be made in this direction. Additionally, our results can be extended to a multivariate analysis aiming at investigating the lead/lag structure of different cryptocurrencies observed at high frequency.

Author Contributions: L.C. conceived of the presented idea. M.S. implemented the analysis under the supervision of L.C. Both authors discussed the results and contributed to the final manuscript.

References

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens. 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics* 61: 43–76. [CrossRef]

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2001b. The distribution of realized exchange rate volatility. *SSRN Electronic Journal* 27: 42–55.

Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold. 2007. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics* 89: 701–20. [CrossRef]

Ardia, David, Keven Bluteau, Kris Boudt, and Leopoldo Catania. 2018. Forecasting risk with markov-switching garch models: A large-scale performance study. *International Journal of Forecasting* 34: 733–47. [CrossRef]

Ardia, David, Keven Bluteau, and Maxime Rüede. 2018. Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters*. [CrossRef]

Balcilar, Mehmet, Elie Bouri, Rangan Gupta, and David Roubaud. 2017. Can volume predict bitcoin returns and volatility? A quantiles-based approach. *Economic Modelling* 64: 74–81. [CrossRef]

Bariviera, Aurelio F. 2017. The inefficiency of bitcoin revisited: A dynamic approach. *Economics Letters* 161: 1–4. [CrossRef]

Barndorff-Nielsen, Ole, Peter R. Hansen, Asger Lunde, and Neil Shephard. 2009. Realized kernels in practice: trades and quotes. *Econometrics Journal* 12: C1–C32. [CrossRef]

Barndorff-Nielsen, Ole E. 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2: 1–37. [CrossRef]

Barndorff-Nielsen, Ole E., and Neil Shephard. 2003. Realized power variation and stochastic volatility models. *Bernoulli* 9: 243–65. [CrossRef]

Barndorff-Nielsen, Ole E., Svend Erik Graversen, Jean Jacod, Mark Podolskij, and Neil Shephard. 2006. A central limit theorem for realised power and bipower variations of continuous semimartingales. In *From Stochastic Calculus to Mathematical Finance*. Berlin/Heidelberg: Springer, pp. 33–68.

Bauwens, Luc, Christian M. Hafner, and Sébastien Laurent. 2012. *Handbook of Volatility Models and Their Applications*. New York: John Wiley & Sons, vol. 3.

Black, Fischer. 1976. Studies of Stock Price Volatility Changes. In *Proceedings of the 1976 Meeting of the Business and Economic Statistics Section*. Washington: American Statistical Association, pp. 177–181.

Bloomberg. 2017. Nasdaq Plans to Introduce Bitcoin Futures. Available online: https://www.bloomberg.com/news/articles/2017-11-29/nasdaq-is-said-to-plan-bitcoin-futures-joining-biggest-rivals (accessed on 10 February 2019).

Breedon, Francis, and Angelo Ranaldo. 2013. Intraday patterns in FX returns and order flow. *Journal of Money, Credit and Banking* 45: 953–65. [CrossRef]

Brownlees, Christian T., and Giampiero M. Gallo. 2006. Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational Statistics & Data Analysis* 51: 2232–245.

Catania, Leopoldo, and Stefano Grassi. 2017. Modelling crypto-currencies financial time-series. *SSRN Electronic Journal*. [CrossRef]

Catania, Leopoldo, Stefano Grassi, and Francesco Ravazzolo. 2019. Forecasting cryptocurrencies under model and parameter instability. *International Journal of Forecasting* 35: 485–501. [CrossRef]

Catania, Leopoldo, Stefano Grassi, and Francesco Ravazzolo. 2018. Predicting the volatility of cryptocurrency time–series. In *Mathematical and Statistical Methods for Actuarial Sciences and Finance: MAF 2018*. Edited by Marco Corazza, María Durbán, Aurea Grané, Cira Perna and Marilena Sibillo. New York: Springer.

CME. 2017. CME Group Announces Launch of Bitcoin Futures. Available online: http://www.cmegroup.com/media-room/press-releases/2017/10/31/cme_group_announceslaunchofbitcoinfutures.html (accessed on 10 February 2019).

Corsi, Fulvio. 2008. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7: 174–96. [CrossRef]

Corsi, Fulvio, Francesco Audrino, and Roberto Renò. 2012. HAR modeling for realized volatility forecasting. In *Handbook of Volatility Models and Their Applications*. Hoboken: John Wiley & Sons, Inc., pp. 363–82.

Creal, Drew, Siem Jan Koopman, and André Lucas. 2013. Generalized Autoregressive Score Models with Applications. *Journal of Applied Econometrics* 28: 777–95. [CrossRef]

Cryptocoinsnews. 2017. Tokyo Financial Exchange Plans for Bitcoin Futures Launch. Available online: https://www.cryptocoinsnews.com/breaking-tokyo-financial-exchange-plans-bitcoin-futures-launch/ (accessed on 10 February 2019).

Dacorogna, Michael M., Ulrich A. Müller, Robert J. Nagler, Richard B. Olsen, and Olivier V. Pictet. 1993. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. *Journal of International Money and Finance* 12: 413–38. [CrossRef]

Diebold, Francis, and Robert Mariano. 1994. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* 13: 253–65.

Dyhrberg, Anne H. 2016. Bitcoin, gold and the dollar—A GARCH volatility analysis. *Finance Research Letters* 16: 85–92. [CrossRef]

Gencay, Ramazan, Michel Dacorogna, Ulrich A. Muller, Richard Olsen, and Olivier Pictet. 2001. *An Introduction to High-Frequency Finance*. Amsterdam: Elsevier Science Publishing Co. Inc.

Hansen, Peter R., and Asger Lunde. 2005. A forecast comparison of volatility models: Does anything beat a garch (1, 1)? *Journal of Applied Econometrics* 20: 873–89. [CrossRef]

Harvey, Andrew C. 2013. *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge: Cambridge University Press, Vol. 52.

Katsiampa, Paraskevi. 2017. Volatility estimation for bitcoin: A comparison of GARCH models. *Economics Letters* 158: 3–6. [CrossRef]

Liu, Lily, Andrew Patton, and Kevin Sheppard. 2015. Does anything beat 5-minute rv? A comparison of realized measures across multiple asset classes. *Journal of Econometrics* 187: 293–311. [CrossRef]

Marcellino, Massimiliano, James H. Stock, and Mark W. Watson. 2006. A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of econometrics* 135: 499–526. [CrossRef]

Nakamoto, Satoshi. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. Available online: https://bitcoin.org/bitcoin.pdf (accessed on 10 February 2019).

Nelson, Daniel B. 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society* 59: 347–70. [CrossRef]

Phillip, Andrew, Jennifer S. K. Chan, and Shelton Peiris. 2018. A new look at cryptocurrencies. *Economics Letters* 163: 6–9. [CrossRef]

Stavroyiannis, Stavros. 2018. Value-at-risk and related measures for the bitcoin. *The Journal of Risk Finance* 19: 127–136. [CrossRef]

Taylor, Stephen J., and Xinzhong Xu. 1997. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4: 317–40. [CrossRef]

Zakoian, Jean-Michel. 1994. Threshold Heteroskedastic Models. *Journal of Economic Dynamics and Control* 18: 931–55. [CrossRef]

Zhang, Yuanyuan, Stephen Chan, Jeffrey Chu, and Saralees Nadarajah. 2019. Stylised facts for high frequency cryptocurrency data. *Physica A: Statistical Mechanics and Its Applications* 513: 598–612. [CrossRef]