Instantaneous Volatility Seasonality of Bitcoin in Directional-Change Intrinsic Time

V. Petrov*†, A. Golub‡ and R. Olsen§

†University of Zurich, 15 Andreasstrasse, 8050 Zurich, Switzerland
‡flov technologies AG, 26 Gotthardstrasse, 6300 Zug, Switzerland
§Lykke Corp., 2 Baarerstrasse, 6300 Zug, Switzerland

We uncover weekly instantaneous volatility seasonality patterns of Bitcoin prices and compare it to the seasonality of three Forex exchange rates. New instantaneous volatility measure based on the analysis of alternating price trends is proposed. The measure is agnostic to the flow of physical time and is equally applicable to all high-frequency tick-by-tick prices. Alternating drawdowns and drawups have been adapted to perform as indicators of the market's activity. We re-expressed these quantities in terms of directional-change intrinsic time. This concept dissects a price curve into a collection of consecutive positive and negative returns measured from local extremes. Interdependence of the instantaneous volatility and the number of directional changes was employed to discover the seasonality patterns. Analytical expressions connecting the waiting times and the number of directional changes have been validated by Monte Carlo simulation. Time intervals throughout which the statistical weekly volatility activity is constant were applied to demonstrate the long memory of the number of directional changes and instantaneous volatility. Provided volatility estimation method can be adapted as a universal multiscale risk-management tool not restricted by the continuity of physical time.

Keywords: Instantaneous volatility; Intrinsic time; Seasonality; Directional-change; Forex; Bitcoin; Scaling laws, Risk management

JEL Classification: G17; G32; C02

Funding details

This project has received funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skodowska-Curie grant agreement No 675044.

*Corresponding author. Email: vladimir.petrov@uzh.ch

1. Introduction

One of the most prominent and still insufficiently studied objects is the rapidly evolving world of cryptocurrencies. The most famous representative of that world is its first successful pioneer Bitcoin. Bitcoin was created in 2008 as an alternative to the classical financial system and as the 'peer-to-peer version of electronic cash' (Nakamoto 2008). Bitcoin and its underlying technology blockchain[1] swiftly gained attention from the technologically savvy community and media and soon became one of the most debatable topics at all levels of the modern society. Over a thousand alternative cryptocurrencies, based on the similar concept, emerged since the time Bitcoin was invented. Some part of them became accessible for trading at various electronic venues also known as crypto-exchanges[2]. In contrast to the traditional FX market, trades of cryptocurrencies happen 24 hours a day and seven days per week. Additionally, there are not many big centralised financial institutions capable to directly influence the state of the system. This happens due to the limited acceptance of this new financial domain by international organisations with the access to sizable funds[3]. These specific facts are the reasons why the seasonality patterns prevalent in the world of cryptocurrencies could be incomparable with the ones typical for the FX market.

Bitcoin's trends drastic changes and their persistence do not only indicate the aggregated expectations and actual actions of all market participants but also reveal its sensitivity to exogenous stress factors. The latter is known as liquidity and is directly related to the scale of market crashes as well as to the periods of recovery (Gennotte and Leland 1990). Asset managers and options seller try to foresee future risks associated with the trend changes by employing a diverse range of risk management tools capable for estimating the probability of price disturbances.

One of the most well-known risk-factors is the probability of so-called drawups and drawdowns, that is, of price drops and price rises between the running price maxima and running minima respectively. Numerous research works have focused on the analysis of the size, periodicity and the time of recovery associated with drawups and drawdowns in traditional markets. The joint Laplace transform was utilised by Taylor (1975) for deriving the waiting time τ_a for a drifted Brownian motion. The joint probability of observing a drawup of a given size after a drawdown, during a given term, was analysed as a homogeneous diffusion process (Pospisil *et al.* 2009). (Zhang 2015) derived the same probability in the context of exponential time horizons (the horizons are exponentially distributed random variables) and described the law of occupation times for both drawup and drawdown processes. These and others theoretical findings were successfully applied to real financial problems briefly discussed in Appendix A.

The cornerstone element of all research works on drawdowns is the first passage time τ_a. This estimate, associated with market crashes of various scales, is especially useful of handling the dynamics of high-frequency markets. The reason for it is the fact that in reality, extreme price jumps occur more often than it should happen when the distribution of returns coincides with the normal one. Fat-tails effect were discovered in the stock market (Jondeau and Rockinger 2003, Rachev *et al.* 2005), in the Forex (Dacorogna *et al.* 2001, Cotter 2005) as well as in Bitcoin prices (Liu *et al.* 2017). Therefore, the efficient set of forecasting techniques aimed at identifying appropriate conditions for the future market crashes should inevitably be supplied by risk-management tools managing sequence of drawdowns and drawups.

Existing literature on risk-management techniques primarily relies on physical time as a measure of the length and periodicity of financial events. In other words, the existence of a universal clock dictating the evolution of the market is assumed. This assumption results in the stochastic volatility of historical returns computed at different scales (Müller *et al.* 1997). More robust techniques which are beyond the limits of physical time applied for studying financial activity are needed to handle this stochasticity. One of the methods capable of doing so is the concept of directional-change

[1] A growing list of records containing information on the ownership of all existing Bitcoins.
[2] Information on all cryptocurrencies and trading venues can be found at the web-resource coinmarketcap.com.
[3] At the moment of writing the paper Wall Street and other big financial hubs are considering trading cryptocurrencies.

intrinsic time (Guillaume *et al.* 1997). This is an event-based framework where the activity of market prices dictates the speed of the transition between different states. The universal physical clock is replaced with an intrinsic one. The method dissects a price curve into sections characterised by alternating trends of arbitrary defined size. The concept is closely related to the evolution of drawdowns and drawups. One can interpret their alternating sequence as a collection of directional changes following each other. Physical time does not play any role in the dissection procedure. More details on the directional-change intrinsic time can be found in Appendix B.

In this research work, we investigate the connection between the observed number of alternating drawdowns and drawups (directional-changes) and the instantaneous volatility. Non-parametric estimation of instantaneous volatility is still a relatively new topic which, to the extent of our knowledge, has not been studied before from the point of view of the directional-change intrinsic time. Obtained analytical expressions are employed to reveal the seasonality structure of Bitcoin's instantaneous volatility. The latter is compared to the seasonality patterns of exchange rates of the traditional FX market. Described experiments contribute to the collection of existing literature on the seasonality properties of Bitcoin and other cryptocurrencies a brief description of which can be found in Appendix C.

The outline of the remaining paper is as follows. Section 2 describes the data used in the experiments and Section 3 outlines the way in which the number of directional changes is connected to the instantaneous volatility. In Section 4 we present all results obtained by the traditional as well as by the novel volatility measurement techniques and also describe the theta time applied to remove the seasonality pattern. Section 5 concludes the main body of the paper and proposes the potential use of the developed technique. Appendix A provides a brief overview of research works on the properties of drawdowns and drawups. Appendix B gives detailed reasoning of the need in the directional-change intrinsic time and contains a set of literature where the concept was successfully applied. In Appendix C the main findings of research works on seasonality patterns of cryptocurrencies are provided. Appendix E includes additional results of the volatility computed by the 'natural' estimator and Appendix D contains results of the Monte Carlo simulation confirming the accuracy of the model.

The fully functional code used in the project can be downloaded from the author's GitHub repository[1].

2. Data

Bitcoin price changes observed at the *Kraken* crypto-exchange were downloaded from the *Bitcoincharts* online platform supplying financial and technical data related to Bitcoin network[2]. The studied time interval is from January 2014 to April 2018 and includes 4 778 429 ticks.

Three FX exchange rates were used in the work to compare their properties with the once found for Bitcoin returns: EUR/USD, EUR/JPY, and EUR/GBP. The covered time interval is from January 2011 to January 2016 and includes 109 069 357, 134 737 397, and 88 704 676 ticks correspondingly. The source of the data is *JForex* trading platform developed by the Swiss bank and marketplace *Ducascopy* which provide various types of the market data in the highest resolution[3].

2.1. *Inner price*

Any collection of historical prices typically assumes two values: the best bid (buy) and the best ask (sell) prices. Non-zero difference between them, called spread, indicates the level of markets' liquidity (Bessembinder 1994, Menkhoff *et al.* 2012). It also has direct connection to realized volatility

[1] shorturl.at/amDLM
[2] http://api.bitcoincharts.com/v1/csv/
[3] https://www.dukascopy.com/swiss/english/forex/jforex/

(Bollerslev and Melvin 1994), and is an indicator of the transaction cost (Hartmann 1999). Another role of the spread is to show the extent of uncertainty the market has on the fair price of the traded asset. The level of the uncertainty constantly changes over time together with the size of the spread. This fact does not allow us to employ only bid or ask prices to study properties of the whole market without missing some part of the information. The average of these two values (mid-price) is also not the best alternative since it does not keep the knowledge on the size of the current spread. Therefore, an alternative measure should be chosen to apply the directional-change algorithm to real data. The concept of inner price was selected to resolve the issue. Inner price can temporarily be equal to the bid or the ask price depending on the direction of the current trend. According to the directional-change algorithm, one should wait for the price increase by δ percents from the local minimum to register a new directional change if the current trend is down. In this case, the value of the local extreme coincides with the best price on the offer side of the order book, that is, the ask price. A new intrinsic event will tick only when the distance between the latest bid price and the local extreme reaches the size of the chosen directional-change threshold δ. Alternatively, the local extreme takes the value of the best bid price and the distance is measured between the newest ask and the extreme if the current mode is upward.

3. Methods

Theoretical researchers mostly rely on the Brownian motion as the proxy for price returns of real markets (the most famous example is the work of Black and Scholes (1973)). The motive behind this is the analogy between historical price moves and changing coordinates of an ensemble of molecules in thermodynamics. In the classical work Osborne (1959) the author shows that the steady-state distribution of log-returns in the stock market is the probability distribution for a particle in Brownian motion. It is important to emphasise that at the time of the publication (1959) the structure and the dynamic of the market was very different from the ones typical for our modern digital world. Instead of stone-made trading floors where all deals happened more than a half a century ago, the present-day trading has almost completely moved to the online space (see Harris (2003) for the historical endeavour on the evolution of trading and exchanges). In this space, a signal can easily propagate through borders with the speed of light and the majority of trades happen in an automated way. The evolution of price returns acquired specific dynamic characterised by a set of stylized facts. These facts, dependent on the selected for the observation time scale, reveal deviations of real returns from the classical Brownian motion (see Cont (2001) for the set of stylized empirical facts). Nevertheless, in our work, the chose of Brownian motion is justified by two reasons. First, the directional-change concept is agnostic to the flow of physical time and is capable of revealing even weak signals hidden in a collection of prices at multiple scales. Second, the divergence of empirical results from the properties of the selected model helps to understand the features of the real world markets better. Therefore, we model the set of prices $\{S_t \colon t \geq 0\}$ as an arithmetic Brownian motion with trend μ and volatility σ:

$$dS_t = \mu dt + \sigma dB_t. \tag{1}$$

In terms of the directional-change intrinsic time framework, $T_{up}(\delta_{up})$ denotes the time for an upward directional change of the size $\delta_{up} > 0$ to unfold. In other words, it is the time interval which passes until the price increases by δ_{up} percents from the local minimum m_t. Technically:

$$T_{up}(\delta_{up}) = \inf\{t > 0 : S_t - m_t \geq \delta_{up}\}, \tag{2}$$

where

$$m_t := \inf_{\epsilon \in [0,t]} S_\epsilon. \tag{3}$$

Similarly, $T_{down}(\delta_{down})$ is the time of a downward directional change of the size $\delta_{down} > 0$:

$$T_{down}(\delta_{down}) = \inf\{t > 0 : M_t - S_t \geq \delta_{down}\}, \tag{4}$$

where

$$M_t := \sup_{\epsilon \in [0,t]} P_\epsilon. \tag{5}$$

Both of these equations are also known in the literature as waiting times of drawups and drawdowns definition of which is provided by equation A3. In Landriault *et al.* (2015) it is shown that expected times of a drawup δ_{up} and a drawdown δ_{down} depend on the volatility and the trend of the drifted Brownian motion and can be mathematically expressed as

$$\mathbb{E}[T_{up}(\delta_{up})] = \frac{e^{-\frac{2\mu}{\sigma^2}\delta_{up}} + \frac{2\mu}{\sigma^2}\delta_{up} - 1}{\frac{2\mu^2}{\sigma^2}}, \tag{6}$$

and

$$\mathbb{E}[T_{down}(\delta_{down})] = \frac{e^{\frac{2\mu}{\sigma^2}\delta_{down}} - \frac{2\mu}{\sigma^2}\delta_{down} - 1}{\frac{2\mu^2}{\sigma^2}}. \tag{7}$$

Using the Taylor expansion $e^{\pm\frac{2\mu}{\sigma^2}\delta} = 1 \pm \frac{2\mu}{\sigma^2}\delta + \frac{(\frac{2\mu}{\sigma^2}\delta)^2}{2!} + \mathcal{O}(\mu^3)$ and letting $\mu \to 0$, one can recover that in the case with no trend the equation simplifies to

$$\mathbb{E}[T_{up}(\delta)] = \mathbb{E}[T_{down}(\delta)] = \frac{\delta^2}{\sigma^2}. \tag{8}$$

These equations establish a scaling law dependence between waiting times of a directional change, instantaneous volatility, and the size of the directional-change threshold. In the analysis of Glattfelder *et al.* (2011) it was empirically found that in the FX market the average expected time is proportional to the size of the directional change threshold δ used to identify alternating trends raised to the power of two:

$$\langle T(\delta) \rangle \sim \delta^2, \tag{9}$$

which confirms the assumption that the evolution of real prices has similar properties to the random walk.

Let $N(\delta_{down}; \sigma, \mu, [0, T])$ denote the number of drawdowns of the size δ_{down} observed within the time interval $[0, T]$ in Brownian motion with parameters μ and σ. Since the sequence $T_{down}(\delta_{down})^1$, $T_{down}(\delta_{down})^2$,... is the sequence of non-negative, independent and identically distributed random variables, the sequence $\{\psi_n; n \in \mathbb{N}\}$ where $\psi_n = T_{down}(\delta_{down})^1 + \ldots + T_{down}(\delta_{down})^n$ is the renewal point process. Thus, $N(\delta_{down}; \sigma, \mu, [0, T])$ can be considered as the renewal counting process and its values can be found using the waiting time equation 7 and applying the Theorem 6.1.1 of Rolski *et al.* (2009) (Landriault *et al.* 2015):

$$\lim_{T\to\infty} N(\delta_{down}; \sigma, \mu, [0, T])^{-1} = \mathbb{E}[T_{down}(\delta_{down})]^{-1} = \frac{\frac{2\mu^2}{\sigma^2}}{e^{\frac{2\mu}{\sigma^2}\delta_{down}} - \frac{2\mu}{\sigma^2}\delta_{down} - 1}. \tag{10}$$

Correspondingly, the expected number of drawups $N(\delta_{up}; \sigma, \mu, [0,T])$ takes the form

$$\lim_{T\to\infty} N(\delta_{up}; \sigma, \mu, [0,T])^{-1} = \mathbb{E}[T_{up}(\delta_{up})]^{-1} = \frac{\frac{2\mu^2}{\sigma^2}}{e^{-\frac{2\mu}{\sigma^2}\delta_{up}} + \frac{2\mu}{\sigma^2}\delta_{up} - 1}. \tag{11}$$

Equations 10 and 11, combined together, give the estimate of the number of directional changes consequently following each other:

$$\mathbb{E}[N(\delta_{up}, \delta_{down}; \mu, \sigma, [0,T])] = \frac{2T\frac{2\mu^2}{\sigma^2}}{e^{-\frac{2\mu}{\sigma^2}\delta_{up}} + e^{\frac{2\mu}{\sigma^2}\delta_{down}} + \frac{2\mu}{\sigma^2}(\delta_{up} - \delta_{down}) - 2}. \tag{12}$$

In the trend-less case the expression is simplified to the following form:

$$\mathbb{E}[N(\delta_{up}, \delta_{down}; \sigma, [0,T])] = \frac{2T\sigma^2}{\delta_{up}^2 + \delta_{down}^2}. \tag{13}$$

These theoretical dependencies between the number of directional changes and the properties of the analyzed time series are equivalent to the empirical observations found in Guillaume *et al.* (1997). There the authors mention that $N(\delta) \sim \delta^{-2}$ (for $\delta = \delta_{up} = \delta_{down}$).

A Monte Carlo statistical test was performed to numerically verify the accuracy of the obtained results expressed in equations 6, 7 and 12. Results of the experiment are provided in Table D1. We selected only positive values of μ since the equations are symmetrical to the direction of the trend. All experiments exhibit high similarity of both empirical and theoretical values.

The meaning behind the provided equations is that the absolute size and the ratio of directional-change thresholds used to dissect a price curve in a sequence of upward and downward trends affect the frequency and the total number of events within a given time interval. In figure 1 we demonstrate three heatmaps where each point corresponds to the number of directional changes registered with a pair of thresholds $\{\delta_{up}, \delta_{down}\}$. Each heatmap represents the results of a Monte Carlo simulation where Brownian motions with different parameters are applied. From equations 10 and 11 it follows that the combination $\gamma = \frac{\mu}{\sigma^2}$ is the crucial factor affecting the expected number of intrinsic events. The expression γ often appears in the intrinsic time framework and is known in the insurance industry as 'adjustment coefficient' or 'the Lundberg exponent' (Asmussen and Albrecher 2010). It finds its application in the ruin theory dating back to 1909 (Lundberg 1909). It is also described as the optimal information theoretical betting size called Kelly Criterion (Kelly Jr 2011). In this work, we check three different scenarios affecting the size of γ: $\frac{\mu}{\sigma^2} = 0$ (figure 1(a)), $\frac{\mu}{\sigma^2} \ll 0$ (figure 1(b)), and $\frac{\mu}{\sigma^2} \gg 0$ (figure 1(c)).

Panel 1(a) in figure 1 corresponds to the set of experiments where the trend of the Brownian Motion is equal to zero. From equation 13 it follows that in such conditions the value $\mathbb{E}[N(\delta_{up}, \delta_{down}; \sigma, [0,T])]$ should be constant along circular contours $\delta_{up}^2 + \delta_{down}^2 = \delta^2$ for $\delta > 0$. The provided picture confirms the noted dependence. In insets 1(b) and 1(c) of figure 1 it is shown that when the 'adjustment coefficient' $\gamma = \frac{\mu}{\sigma^2}$ is significantly smaller or bigger than zero, the circular contours transform into ellipses. This phenomenon can be interpreted in the following way: if $\mathbb{E}[N(\delta_{up} = \delta_{down}; \gamma = 0, [0,T])]$ is the expected number of directional changes registered in the drift-less time series of given length and fixed σ then for any γ greater or smaller then zero there is always such a couple of non equal thresholds $\{\delta_{up}, \delta_{down} | \delta_{up} \neq \delta_{down})\}$ that

$$\mathbb{E}[N(\delta_{up}, \delta_{down} \mid \delta_{up} \neq \delta_{down}; \gamma \neq 0, [0,T])] = \mathbb{E}[N(\delta_{up}, \delta_{down} \mid \delta_{up} = \delta_{down}; \gamma = 0, [0,T])]. \tag{14}$$

The property is essential for risk management techniques constructed on top of directional-change intrinsic time: any process characterized by certain degree of persistent trend could be treated as the one without the trend by tuning the size and the ratio of selected directional-change thresholds.

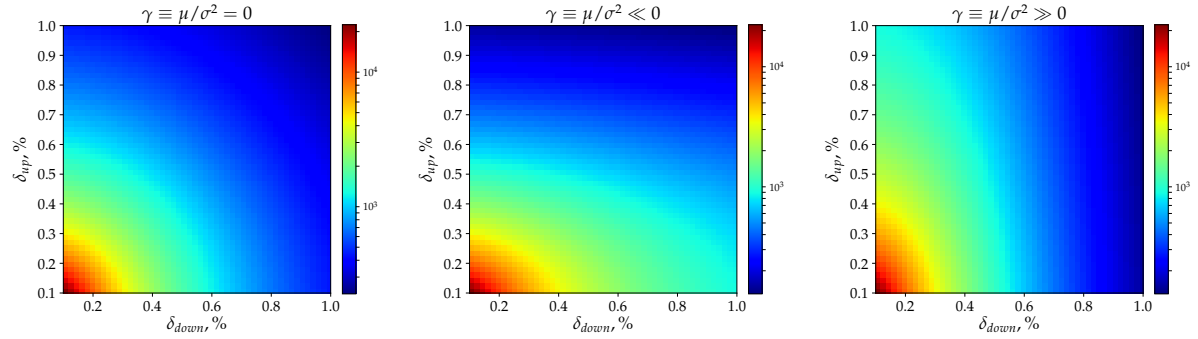

(a) (b) (c)

Figure 1. Heatmap of the number of directional changes. Values coincide with the ones computed using equation (13) and various combinations of directional-change thresholds $(\delta_{up}, \delta_{down})$ and parameters of Brownian motion (μ, σ).

An example of real application of this fact is provided in Golub *et al.* (2017) where the authors employ it to design optimization inventory control function sensitive to the trend changes.

The clear understanding of the way how volatility changes over time is particularly important for risk management and inventory control problems. As discussed in Section B, classical volatility estimation methods, also called 'natural' estimators (Cho and Frees 1988)[1], primarily rely on physical time as the core reference for historical returns. Stochastic volatility accounted for this fact became a cornerstone for multiple research works (for example, Andersen and Lund (1997), Barndorff-Nielsen and Shephard (2002), Aı *et al.* (2007) and many others). Values, computed by 'natural' estimators, dominantly correspond to the integrated volatility of the studied process. Alternative estimators, designed to reveal the size of the instantaneous volatility, mostly based on Fourier analysis and require extensive computations (see Chapter 3 in Mancino *et al.* (2017)).The directional-change concept is by design agnostic to the speed of the time flow and it automatically adapts to the periods of changing activity. This property of the directional-change intrinsic time together with analytical equations 12 and 13 bring the idea of a new instantaneous volatility estimator devoid of the equidistant timestamps shortcomings. From equation 13 it follows that for a trend-less time series the instantaneous volatility can be estimated by counting the number of directional-changes within the time interval $[0, T]$:

$$\sigma_{DC} = \delta \sqrt{\frac{N(\delta)}{T}}. \tag{15}$$

We put the superscript DC to distinguish the volatility computed through the directional-change intrinsic time from the traditional estimators. The latter we will mark by σ_{trad}.

Equation 15 solely computes the volatility part of the process in case of a random walk. Therefore, it can be classified as the true estimator of the instantaneous volatility. In the current work, we apply equation 15 to study changing dynamic of financial time series throughout one week to reveal its seasonality pattern.

[1]The work Cho and Frees (1988) is particularly interesting due to the analysis the authors did to compare volatilities computed by 'natural' and 'temporal' estimators. The latter employs time intervals measured between consequent and alternating price moves of fixed relative size and thus is very close to the approach presented in the current paper.

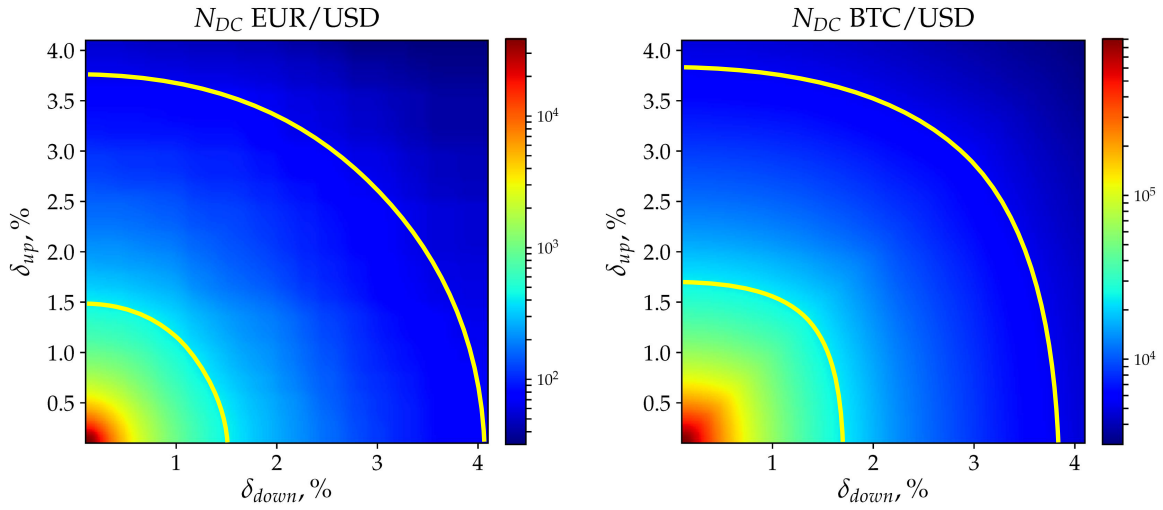

(a) (b)

Figure 2. Heat map of the number of directional changes calculated in (a) EUR/USD and (b) BTC/USD time series. Heatmaps have different scales.

4. Results

4.1. *Number of directional changes*

Historical returns from real markets have properties similar to the Brownian motion used to derive equations 12 and 13 connecting the expected number of directional changes with parameters of the underlying process. Similar counters shown in figure 1 should be found in heatmaps depicting the number of directional changes empirically registered in real data if the assumption of the normal distribution of returns is true. EUR/USD and BTC/USD exchange rates were taken to verify the statement by replicating the same experiment done with the Brownian motion before. A collection of 40 directional-change thresholds ranging from 0.1% to 4.1% defines the scale of the heatmap grid. Results of the experiments are presented in figure 2. Colour schemes, used for both plots, have different scales due to the difference in EUR/USD and BTC/USD volatility. Yellow lines indicate areas of the equal number of directional changes corresponding to $\delta = \{1.15\%, 2.8\%\}$ in case of EUR/USD and to $\delta = \{1.4\%, 3.0\%\}$ in case of BTC/USD. Curves in figure 2(a) have almost circular shape and are only slightly shifted towards the bigger δ_{down} values. This shift is present due to the downward trend experienced by the exchange rate from 2011 to 2016 (from \$1.4 to \$1.1 per Euro). EUR/USD time series exhibited relative stability with no noticeable regime switches apart from this constant trend. BTC/USD exchange rate was much more unstable concerning the dynamic of EUR/USD. The price was relatively steady and did not demonstrate any permanent non-zero trend from 2014 to 2017. The price of Bitcoin grew with accelerating pace by more than 20 times in the second half of 2017 and then lost nearly 70% of its value at the beginning of 2018[2]. These significant trend changes are illustrated in figure 2(b) by yellow counters with a notable deviation from the circular shape. This shape can be decomposed into two parts of independent ellipses similar to the ones observed for Brownian motion with non-zero 'adjustment coefficient' γ (Figures 1(b) and 1(c)). The price roller-coaster caused considerable disparity of the number of registered directional changes by any pair of thresholds before and after the pick at the end of 2017.

[2]It had a minimum at \$230 per Bitcoin, temporary maximum at \$20 000, and then a drop to \$6 000.

4.2. *Realized versus instantaneous volatility*

In the second experiment, we compared the annualised volatility computed by the traditional method (equation 16) and the volatility based on the observed number of directional changes (equation 15). Returns R_t are computed as logarithms of the ratio of consecutive prices S_t and S_{t-1} in the traditional technique. The length of the sample n depends on the selected time interval between two measures of the returns. The whole sample is used to find the standard deviation of the time series also knows as realised volatility σ_{trad}:

$$R_t = \ln(S_t/S_{t-1}), \quad R_{avg} = \frac{\sum_{t=1}^{n} R_t}{n}, \quad \sigma_{trad} = \sqrt{\frac{\sum_{t=1}^{n}(R_t - R_{avg})^2}{n-1}}. \tag{16}$$

The directional-change method does not define the number of observations ex-ante in contrast to the traditional approach where the length of a sample depends on the preselected time span between consecutive returns. According to equation 15 the size of the directional-change threshold δ determines the expected range of the number of prices used to calculate the instantaneous volatility. This flexibility of the intrinsic time makes it possible to use the data of the highest frequency: tick-by-tick prices.

Parameters of tools used to estimate volatility can affect the results of the experiment (Müller *et al.* 1997). Therefore, four increasing time intervals Δt_k where $k = \{1, 2, 3, 4\}$ were selected to define the distance between each pair of consecutive prices S_t and S_{t-1} used for the 'natural estimator': $\Delta t_1 = 1$ minute, $\Delta t_2 = 10$ minutes, $\Delta t_3 = 1$ hour, and $\Delta t_4 = 1$ day. The set of thresholds employed to investigate the directional-change approach can also be arbitrarily chosen. However, we initially targeted the goal to compare the results of both experiments. For this reason, we used the number of timestamps corresponding to each time interval Δt_k as the target for the number of directional changes registered in the same data set. That is, the collection of four thresholds δ_k was selected in such a way that in the given time series the number of directional changes will be approximately equal to the number of time intervals n_k of the length Δt_k. We utilised one of the scaling properties described in Glattfelder *et al.* (2011) to find the precise thresholds size. The scaling property has the name 'Time of total-move' scaling law (law 10 in the article) where the total-move is composed of a directional-change (DC) and an overshoot (OS) parts. The law connects the size of the threshold δ with the waiting time $T_{TM}(\delta)$ between two consecutive intrinsic events:

$$T_{TM}(\delta) = \left(\frac{\delta}{C_{t,TM}}\right)^{E_{t,TM}}, \tag{17}$$

The currency average scaling parameters $E_{t,TM}$ and $C_{t,TM}$ computed in Glattfelder *et al.* (2011) are 2.02 and $1.65e{-}3$, correspondingly. Putting these coefficients into equation 17 one can calculate that thresholds reciprocal to the selected time intervals are: $\delta_1 = 0.013\%$, $\delta_2 = 0.039\%$, $\delta_3 = 0.095\%$, and $\delta_4 = 0.458\%$. It is worth mentioning that applied scaling parameters are relevant only to the FX market which was the object of the research in Glattfelder *et al.* (2011). To the extent of our knowledge, parameters specific to Bitcoin prices were not mentioned in the scientific literature before. Therefore, as the first step, we obtained the parameters by studying the 'time of total-move' scaling law of historical Bitcoin returns. In figure 3(a) the log-log plot of waiting times $T_{TM}(\delta)$ is provided. The red line marks BTC/USD scaling law and is shown together with black and green lines computed for EUR/USD and Geometrical Brownian Motion (GBM). Settings of the latter are chosen to mimic realistic returns typical for the FX market.

Scaling law parameters, obtained in the experiment, exhibit distinct resemblance of the stylized properties of the traditional FX and the emerging Bitcoin markets. Scaling factors $E_{t,TM}$ of EUR/USD and BTC/USD are 1.827 and 1.818 correspondingly ($\approx 0.5\%$ difference). The same scaling factor of the GBM is 1.920 ($\approx 5.6\%$ difference with EUR/USD) which is noticeably distant

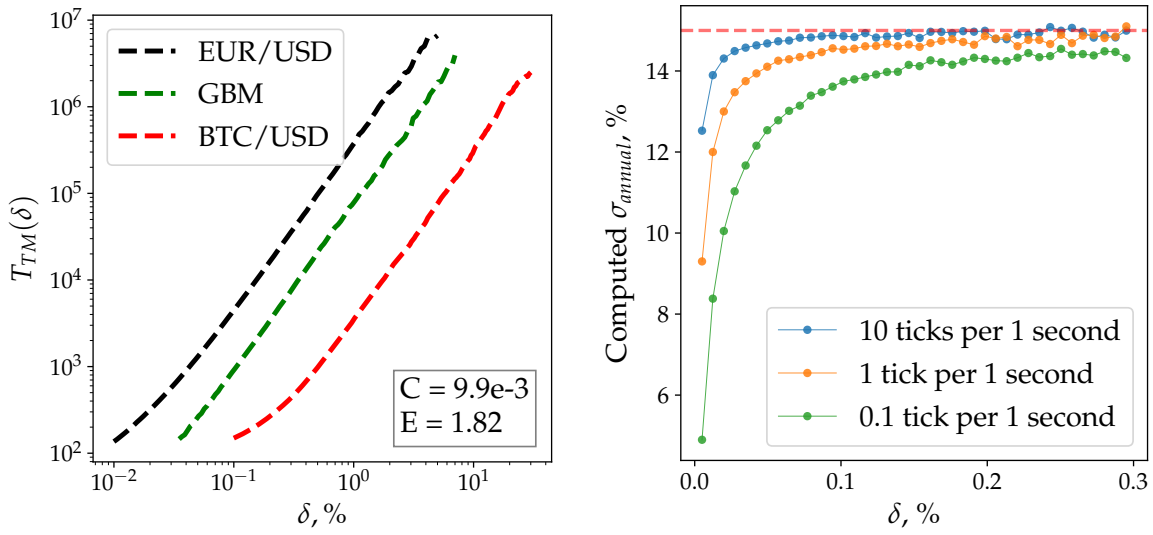

(a) (b)

Figure 3. (a) Total-move scaling laws computed for BTC/USD, EUR/USD, and Geometrical Brownian Motion (GBM). GBM's parameters are $S_0 = 1.3367$, $\mu = 0$, $\sigma = 20\%$, $T = 1$ year, and 10 millions ticks in total. Scaling parameters in the inset are provided for the BTC/USD curve. (b) Instantaneous volatility of three time series generated by GBM with various tick frequencies and fixed volatility (15%) computed by the directional-change approach (equation 15). Sizes of thresholds, used to calculate the values, are put on the x-axis. Red dashed line marks the 15% level.

from the parameters of the analysed exchange rates. This divergence is present due to the non-normal distribution of real returns at ultra-short timescales (fat tails). The effect is pronounced in figure 3(a) as the upward bend of the curves towards the beginning of the x-axis. Linear regressions, built in the range of straight parts of the curves, return scaling coefficients $E_{t,TM}$ of real exchange rates which are very close to the ones observed in GBM. This evidence is an additional confirmation of the 'Aggregational Gaussianity' stylized fact typical for high-frequency markets (Cont 2001). Scaling parameters $C_{t,TM}$ of EUR/USD and BTC/USD are $9.07e-4$ and $28.35e-4$ correspondingly. These values are significantly different because of the unlike scale of the Bitcoin and EUR volatility and are not critical for the analysis.

Scaling law parameters $E_{t,TM}$ and $C_{t,TM}$ found for BTC/USD prices were used to compute the size of directional-change thresholds which would in an average register the number of intrinsic events equal to the number of periods n_k. Expressing the parameter δ_k from the equation 17 we find that for BTC/USD the thresholds are: $\delta_1 = 0.09\%$, $\delta_2 = 0.33\%$, $\delta_3 = 0.89\%$, $\delta_4 = 5.13\%$. The values are about 10 times bigger than the ones for the FX market because of the proportionally larger realized volatility.

The set of selected time intervals and the corresponding thresholds for two markets were applied to compute realized and instantaneous volatility with the traditional and the novel approach. In Table 1 we present: average value of the realized volatility $\langle\sigma_{trad}\rangle$ computed as the sum of all four measurements with $k = \{1, 2, 3, 4\}$ divided by the number of experiments; its standard deviation σ^{-}_{trad}; average value of the instantaneous volatility computed by the novel approach $\langle\sigma_{DC}\rangle$; the corresponding standard deviation σ^{-}_{DC}; ratio of both measures $\langle\sigma_{trad}\rangle/\langle\sigma_{DC}\rangle$ and $\sigma^{-}_{trad}/\sigma^{-}_{DC}$. The last column of the table demonstrates the difference in the stability of results obtained by two measures. The gap between the sizes of the realized and the instantaneous volatility is significant and is pronounced across all tested exchange rates. The volatility computed in the 'natural' way persistently exceeds the instantaneous volatility discovered via the novel approach. The divergence growths up to 1.40 times in case of EUR/JPY. Two types of Bitcoin's volatility appear to be only 5% different (column $\langle\sigma_{trad}\rangle/\langle\sigma_{DC}\rangle$). This phenomenon is captivating especially taking into

Table 1. Volatility computed using the 'traditional' (equation 16) and the directional-change approaches (equation 15). Provided values $\langle\sigma_{trad}\rangle$ and $\langle\sigma_{DC}\rangle$ are the average of four measurements performed with specific parameters: in the 'traditional' case time intervals between observations S_n and S_{n-1} are 1 minute, 10 minutes, 1 hour, and 1 day and in case of the novel approach the thresholds δ are $\delta_1 = 0.013\%$, $\delta_2 = 0.039\%$, $\delta_3 = 0.095\%$, $\delta_4 = 0.458\%$ (FX prices) and $\delta_1 = 0.09\%$, $\delta_2 = 0.33\%$, $\delta_3 = 0.89\%$, $\delta_4 = 5.13\%$ (BTC prices).

Name	$\langle\sigma_{trad}\rangle, \%$	$\bar{\sigma}_{trad}$	$\langle\sigma_{DC}\rangle, \%$	$\bar{\sigma}_{DC}$	$\langle\sigma_{trad}\rangle/\langle\sigma_{DC}\rangle$	$\bar{\sigma}_{trad}/\bar{\sigma}_{DC}$
EUR/USD	9.72	0.03	7.53	1.38	1.29	0.02
EUR/JPY	11.93	0.12	8.55	2.07	1.40	0.06
EUR/GBP	8.04	0.23	5.81	1.43	1.38	0.16
BTC/USD	84.76	8.67	80.87	22.21	1.05	0.39

account that Bitcoin is particularly famous due to its oversized price activity. Its activity is clearly pronounced in the large standard deviation of the BTC/USD pair. FX exchange rates, having noticeably smaller realized volatility, are characterized by the wider range of the standard deviation values. The ratio $\bar{\sigma}_{trad}/\bar{\sigma}_{DC}$ reaches the 0.02 level computed for EUR/USD.

4.3. *Discrete price impact*

The standard deviation computed for four different directional-change thresholds has extremely high value. This indicates that in contrast to the realised volatility the instantaneous one does not scale together with time. The cause of the detected aspect is the price discontinuity typical for all real markets. Conventional exchange architecture restricts the price quotations to be a multiple of some constant, for example, 0.001 of a USD. This fact caused substantial debates in the scientific literature with regards to the accuracy of the 'natural' estimators and on the extent to which they overestimate the actual volatility (Gottlieb and Kalay 1985, French and Roll 1986). Equation 15 is built around the continuous Brownian motion and connects the number of directional changes with the instantaneous volatility. It has no adjustment factors to the discreteness of the analysed process. The directional-change intrinsic time does not precisely tick at the level where the size of the return is **equal** to the size of the threshold δ if the size of a price step is relatively big. Instead, in most cases, a new directional-change event becomes registered when the price has already jumped **over** the expected level (slippage). That is, two factors contribute to the size of the instantaneous volatility computed by the novel approach: the scale of the selected threshold δ and the tick size in the given data sample (discreteness).

We performed the following experiment to estimate the impact of the price discreteness and the threshold size δ on the computed instantaneous volatility. Three time series were generated by GBM with the various density of ticks per period of time. Variation of the number of price changes in the simulation is equivalent to changing the simulated tick size having fixed a one-year time interval and volatility of the generated process equal to 15%. Forty directional-change thresholds with fixed increment starting from 0.01% and ending at 0.29% were applied to all three GBMs. We provide computed instantaneous volatility of simulated time series in figure 3(b). Two particular properties can be noticed. First, one brings the generated time series closer to the continuous process by making the size of a tick smaller (increasing the number of price changes in the sample). In this case, the values of the estimated instantaneous volatility σ_{DC} become closer to the volatility σ embedded in the model. Second, bigger thresholds are less sensitive to the discreetness of the given set of prices. The slippage effect becomes less pronounced and the obtained result also approaches the value σ when the tick size represents a small fraction of the directional-change thresholds. Both effects emphasise the dissimilarity of realised and instantaneous volatilities. A more comprehensive analysis should be performed to bridge the gap between these two measures.

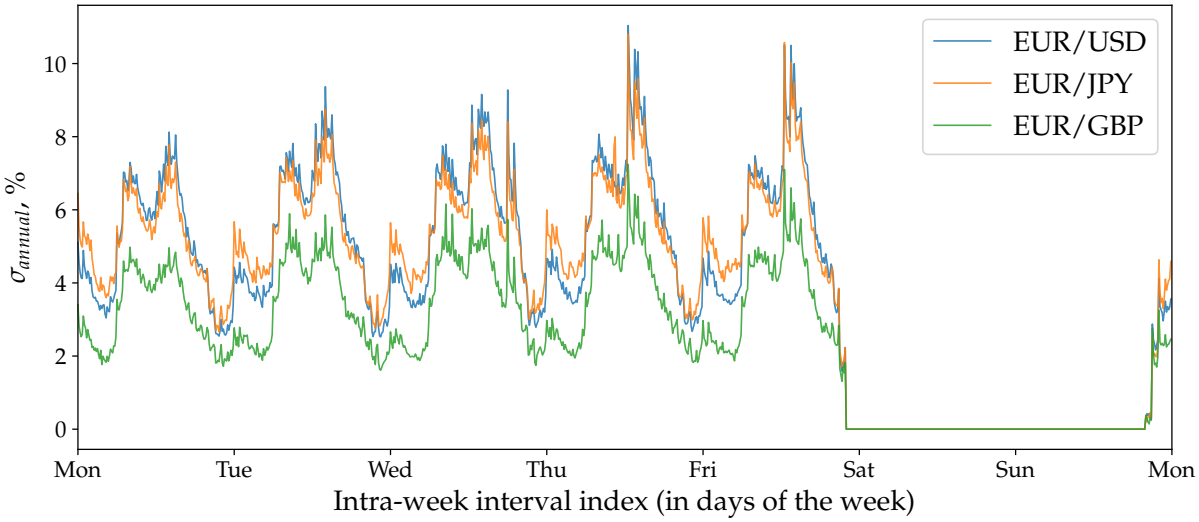

Figure 4. Instantaneous volatility seasonality of three FX exchange rates computed using equation 15. Applied directional-change threshold is $\delta = 0.01\%$. The whole week is devided by equally spaced time intervals $T = 10$ minutes (1008 bins in total).

4.4. *Volatility seasonality*

The seasonality of realised volatility in the traditional markets is a relatively well-studied topic. Dacorogna *et al.* (1993) presented a weekly seasonality pattern of price activity in the FX market. Their analysis depends on the assumption that trading happens at various time zones and within specific trading hours. Such a physical distribution of traders is embodied in geographical components of the market activity. We do not build on a similar assumption in our work. The collection of observed historical returns is treated as the only source of available for the analysis information.

We divide a whole week into a set of 10-minute time intervals (bins) for each of which the average number of directional changes will be computed to find the seasonality pattern of instantaneous volatility typical for Bitcoin prices and FX exchange rates. There are 1008 equally spaced points located on the fixed distance from the beginning of each week in total. This is significantly large number than the one used in the work Dacorogna *et al.* (1993) (168 points). We can afford this decreased granularity thanks to the more detailed historical time series employed for the experiment: instead of 12 million ticks for 26 exchange rates, we have in average 100 million ticks for each FX pair.

For the first experiment with FX exchange rates, we select the threshold $\delta = 0.01\%$. The average number of directional changes in a week registered by a threshold of this size is approximately equal to the number of bins in it. The computation of the seasonality pattern is done in the following steps. First, we run all historical tick-by-tick prices through the directional-change algorithm with the specified threshold δ. As soon as a new intrinsic event becomes registered we check within which bin it happened and add $+1$ to the number of directional changes corresponding to this time interval. When there are no prices left in the historical time series, we find the average number of intrinsic events per each bin and apply equation 15 to compute the corresponding instantaneous volatility. Considering the five-year-long historical data, the obtained average is based on nearly 250 observations. Calculated values should be later normalised by the number of years to get the annualised volatility.

The reconstructed instantaneous volatility seasonality pattern of the FX pairs is shown in figure 4. The pattern is notably stable across all tested exchange rates and is similar to the one demonstrated in Dacorogna *et al.* (1993). This similarity confirms the idea that the seasonality of instantaneous volatility is subject to the geographical distributing of the trading centres across the world.

We provide results of the same experiment where the 'traditional' volatility estimator (equation

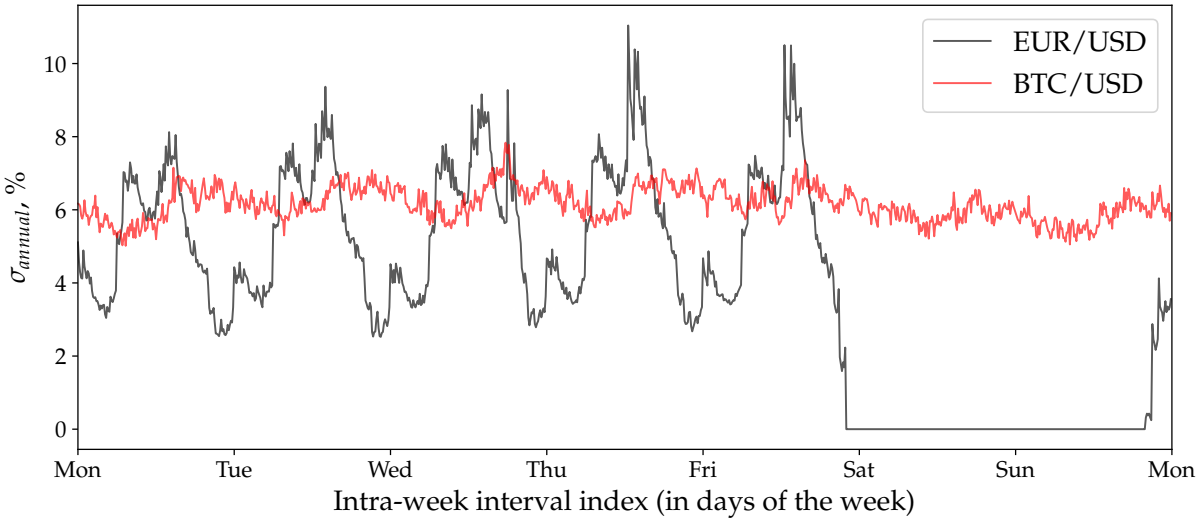

Figure 5. Instantaneous volatility seasonality of BTC/USD compared to the seasonality pattern of EUR/USD computed using equation 15. Directional-change threshold $\delta = 0.01\%$ is used in both experiments. Each time interval $T = 10$ minutes. 1008 bins in total.

16) was employed to reveal the seasonality patterns of the FX exchange rates in Appendix E. The 'traditional' pattern is less affected by the source of the given time series. The average difference between the realized volatility of the most active pair (EUR/JPY) and the least active (EUR/GBP)[1] is equal to 46%. The difference in the volatility computed by the novel approach rises to 56%. The vertical distance between EUR/USD and EUR/GBP curves also becomes less pronounced computed by the 'natural' estimator. At the same time, this estimator demonstrates more rapid changes between values of each couple of consecutive bins. Local maximums at the beginning and the end of a day are considerably abrupt. The reason of this is that the directional-change intrinsic time captures the part of the volatility free of the noise component of the underlying process. The exact form and scale of this part is a topic for the future research work.

Assets traded in the crypto market have several specific properties which make them noticeably distinct from the traditional financial instruments such as FX exchange rates. Among the characteristics are open trading within weekends and holidays; the absence of isolated physical trading centres where working hours are fixed; still low acceptance of the emerging market among professional traders. We explored the seasonality pattern of Bitcoin instantaneous volatility to check whether these specialities have any considerable impact on its shape behaviour. The seasonality pattern is presented in figure 5. We apply the same threshold size $\delta = 0.01\%$ used in the previous experiment to compare the seasonality patterns of Bitcoin and EUR. In contrast to EUR/USD, the periodical shape of Bitcoin's curve is much less pronounced. Its standard deviation computed within a week is 0.5% which is roughly four times smaller than the standard deviation of the EUR/USD pattern equal to 1.9%. Surprisingly, the intra-day maximums and minimums do not precisely coincide with those observed in the traditional markets. They are shifted towards the time intervals where European and American markets contribute the most to the geographical pattern disclosed in Dacorogna *et al.* (1993). This observation confirms the one provided in Eross *et al.* (2017). Realized volatility over weekends is slightly lower than the middle part of the week and is practically equal to the Monday's activity. We attribute the observed facts to the mentioned above non-traditional characteristic of the cryptocurrency market.

The computed instantaneous volatility by the novel approach directly depends on the size of the selected δ (figure 3(b)). To examine the impact of the threshold size we arbitrary selected

[1] According to the Table 1.

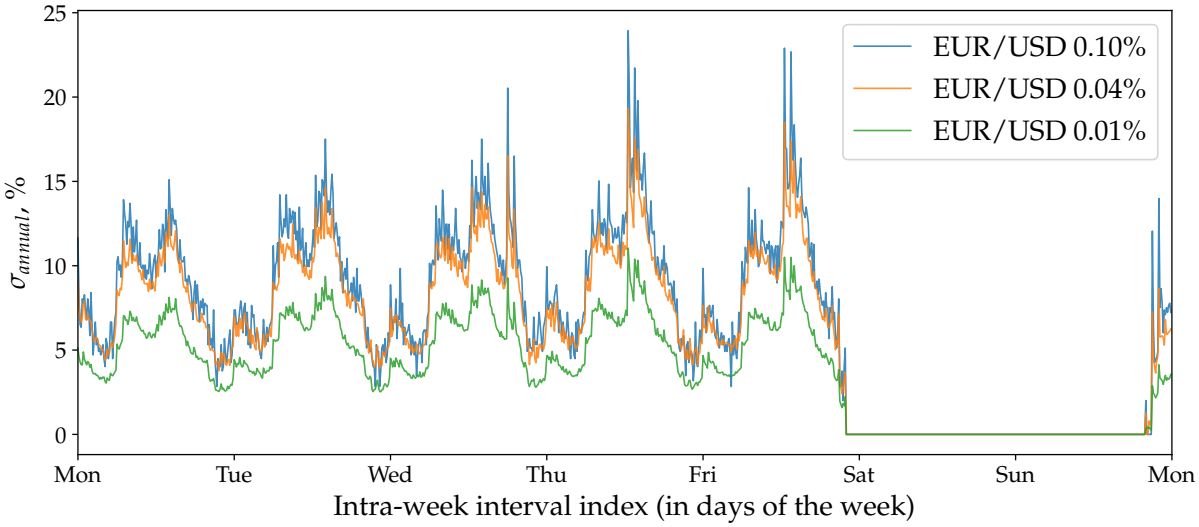

Figure 6. Volatility seasonality of EUR/USD computed using the novel approach (eq. 15) and three different thresholds: $\delta = \{0.01\%, 0.04\%, 0.10\%\}$. The size of a bin is 10 minutes, 1008 bins in a week.

the following set of values $\delta = \{0.01\%, 0.04\%, 0.10\%\}$. The same algorithm described above was applied to reconstruct the volatility seasonality pattern for the FX pair with the highest daily trading volume: EUR/USD (BIS 2016). The seasonality patterns shift toward higher volatility values when the size of the threshold is bigger which is in line with the results of the experiments on GBM (figure 3(b)). Average values of the seasonality curves computed with thresholds δ equal to 0.10% and 0.04% are 1.71 and 1.57 times higher than the values computed with $\delta = 0.01\%$. The difference in the amplitude of all patterns is even more pronounced: the seasonality curve constructed with the smallest in the set threshold is much sleeker (less wander) than the rest of the curves. This phenomenon should urge researchers and practitioners to select threshold according to their needs very carefully employing the directional-change technique.

According to Table 1, realised volatility of Bitcoin returns computed in the 'traditional' way is about nine times bigger than the analogous volatility of the FX exchange rates. Besides, the retrieved sample of historical BTC/USD prices has 1.2 million ticks per year which is 16.7 times smaller than the number of ticks per year in the EUR/USD case (about 20 millions). As a result, the choice of the directional-change threshold δ has a much more significant effect on the average value of the BTC/USD instantaneous volatility. We demonstrate results of four experiments with different threshold sizes in figure 7. The same $\delta = 0.01\%$ is used as the reference for the set of all thresholds: $\delta = \{0.01\%, 0.03\%, 0.10\%, 0.20\%\}$. As it can be seen from figure 7, the increase in the size of δ causes the corresponding increase in the volatility level around which the seasonality pattern oscillates. The levels of the seasonality distribution for $\delta = \{0.03\%, 0.10\%, 0.20\%\}$ are 1.5, 4.0, and 11.1 times bigger than the value corresponding to the smallest threshold $\delta = 0.01\%$. The biggest $\delta = 0.20\%$ lifts the value up to the level of 68.5% (which is still smaller than the realized volatility presented in Table 1 (84.76%)).

4.5. *Autocorrelation and theta-time*

Persistent seasonality patterns of the instantaneous volatility computed for the FX exchange rates have a shape which changes with clear daily periodicity. This observation suggests that there should be a strong autocorrelation of the instantaneous volatility or, in other words, of the number of directional changes. To check this statement, we examined the autocorrelation function (ACF) of the number of directional changes observed within each bin of a week. Results of the experiment made for the FX exchange rates are provided in figure 8. It is not wrong to say that figure 8 also

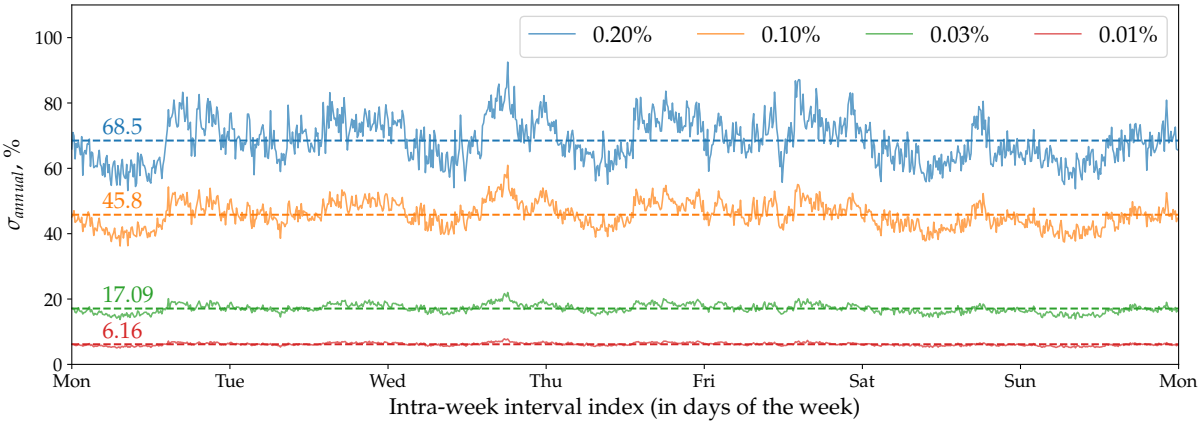

Figure 7. Instantaneous volatility seasonality of BTC/USD returns computed using the directional-change approach (equation 15). Applied thresholds: $\delta = \{0.20\%, 0.10\%, 0.03\%, 0.01\%\}$. Bin size $T = 10$ minutes in all cases (1008 bins in a week). The dashed lines and the numbers over them represent the average level of each seasonality pattern.

indicates the ACF of the volatility since the number of directional changes is directly related to the instantaneous volatility through the equation 15. The same size of the directional-change threshold used to reveal the seasonality distribution $\delta = 0.01\%$ was employed. A remarkably stable pattern was found where daily and weekly seasonality is easily recognisable. Although the ACF function of FX exchange rates discovered in our work is highly similar to results provided by Dacorogna *et al.* (1993), there are clear differences between the two patterns. The ACF of the number of directional changes computed through time lags defined in physical time does not cross the zero level for a much more extended period and is consistently positive with lags even greater than several weeks. The curve representing the ACF of EUR/JPY has the smallest amplitude (smallest variability). In contrast, curves of EUR/USD together with EUR/GBP invariably follow the same pattern shifted up in the case of EUR/USD.

The BTC/USD exchange rate characterized by much less pronounced seasonality pattern of the instantaneous volatility has also been tested by the autocorrelation function. The results are presented in figure 8. The main difference between the values computed for two markets is the amplitude of the ACF curves: the variability of the BTC/USD curve is 10 times lower than the variability of the EUR/USD curve. We note that significantly bigger thresholds have also been tested, but they reveal less accurate patterns due to insufficiently frequent data for the statistical analysis.

It can be seen from figure 8 that the ACF of the FX exchange rates exhibits a certain level of decline. The large seasonal peaks of the autocorrelation functions in physical time do not allow to measure its decay precisely. A measure capable of converting the price evolution process to the stationary one should be applied to better estimate the level of the downturn. We removed the seasonality pattern by employing the concept of theta time (Θ-time) proposed by Dacorogna *et al.* (1993). Θ-time is designed to eliminate the periodicity pattern by defining a set of non-equal intervals within which the measure should be performed. The length of each interval in Θ-time depends on the historical activity of the market rather than on equally spaced periods of the homogeneous physical time. Therefore, the cumulative price activity (volatility) between each consecutive couple of Θ steps is constant. The distance between Θ timestamps, measured in physical time, is dictated by the shape of the volatility seasonality pattern and, in contrast to physical time initially used to reveal the pattern, is not a fixed value. Periods of high activity are equivalent to shrinking the speed of physical time and the frequency of Θ stamps increases. In contrast, periods of low activity are identical to stretching the flow of the physical time and the lower number of Θ intervals appears. Active parts of the seasonality pattern have the higher density

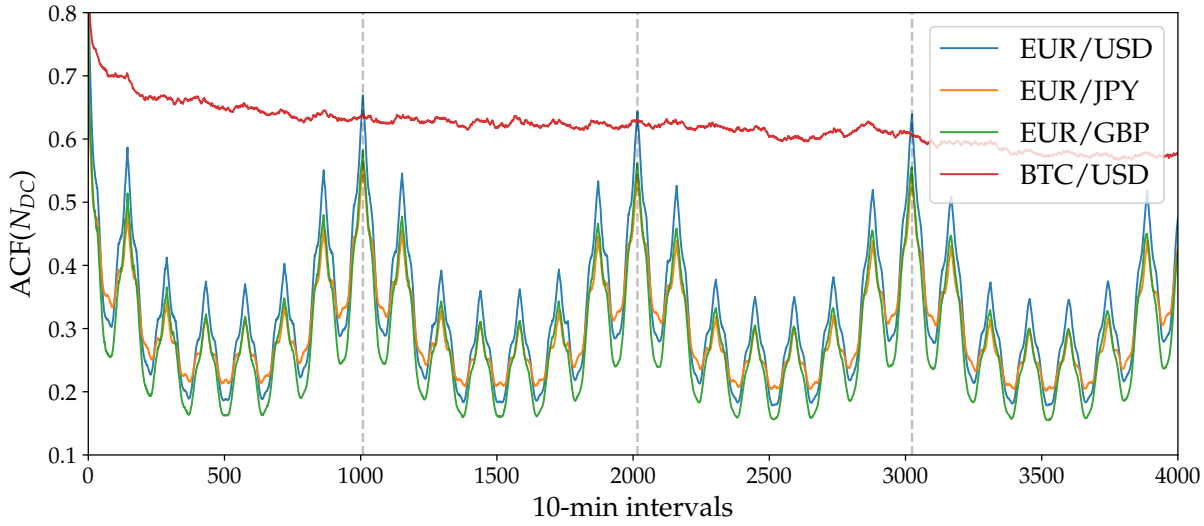

Figure 8. Autocorrelation function of the number of directional changes per a bin in physical time. The bin size $T = 10$ minutes. Vertical dashed lines label weekly intervals. Applied threshold $\delta = 0.01\%$.

of Θ timestamps per a unite of the physical time than the standstill sections. Mathematically:

$$\Theta(t) = \int_{t_0}^{t} \sigma(t')dt', \tag{18}$$

where t_0 and t are the beginning and the end of the considered period of physical time and $\sigma(t')$ is the value of the instantaneous volatility corresponding to each moment of the interval. Equation 18 transforms into the sum of elements $\sigma_{\Delta t'}$ between the beginning and the end of the observed interval Δt_0 and Δt in case of non-continuous seasonality pattern where the values are discretely defined in periods Δt (as in our experiment):

$$\Theta(t) = \sum_{\Delta t_0}^{\Delta t} \sigma_{\Delta t'}. \tag{19}$$

It should be noted that the number of bins in a week is always constant in both physical and Θ times. This is achieved through the assumption that the integral (or the sum) of the weekly activity is the constant value.

The autocorrelation function of the number of directional changes computed in Θ-time is shown in Figures 9(a) and 9(b) in normal and log-log scales. Curves on the log-log plot outline the first 100 bins and are approximated by straight lines corresponding to the scaling law with coefficients E_{ACF} and C_{ACF} provided in Table 2. Major weekly fluctuations of the volatility seasonality patter have been successfully eliminated. Nevertheless, the picture has several peculiar properties which should be discussed. First, Θ-time does not completely remove the seasonality shape of ACF in the same way it happened in the work Dacorogna *et al.* (1993): noticeable peaks are still present in the final part of each business day. The similar phenomenon, observed in the original paper, was explained by the non-optimal setup of the chosen model: it assumed the same activity for all working days, which is, indeed, not fully correct (see figure 4). However, we do not use any analytical expression postulating the equal daily activity to describe the seasonality pattern. Instead, components $\sigma_{\Delta t'}$ of real empirically found volatility seasonality patterns depicted in Figures 4 and 5 were utilized to define the timestamps in Θ-time. Therefore, we eliminate the inefficiency connected to the assumption mentioned above which means an alternative interpretation for the remained seasonality should be provided.

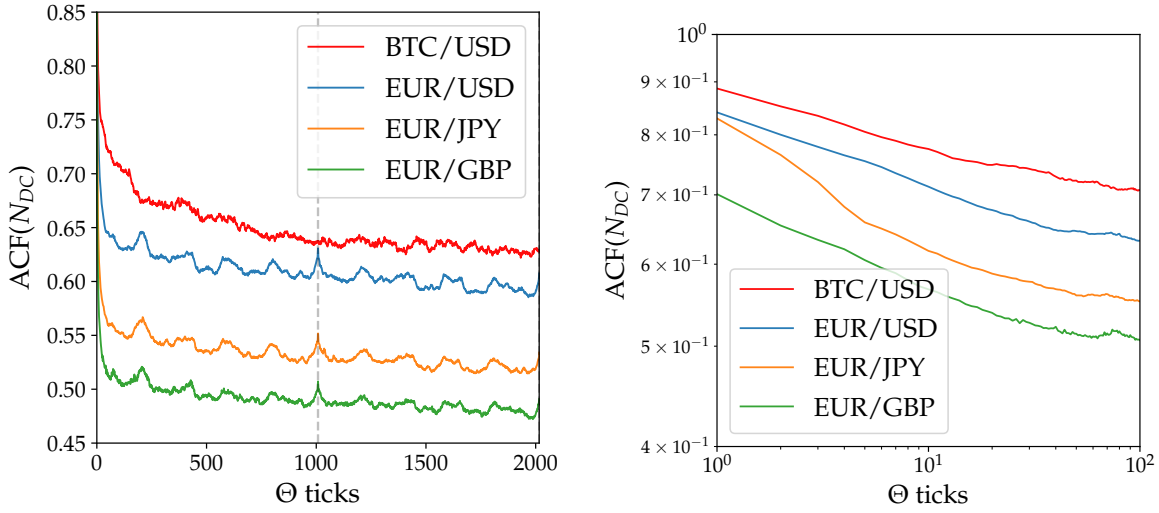

(a) (b)

Figure 9. (a) Autocorrelation function (ACF) of the number of directional changes per a bin in Θ-time. Vertical dashed line labels one week interval. Its location coincides with the location of the vertical line depicted in figure 8. (b) The same ACF on the log-log plot for the beginning of a week.

Table 2. Parameters of the scaling law describing the exponential decay of the ACF of the number of directional changes in Θ-time (figure 9(b)).

Name	C_{ACF}	E_{ACF}	R
BTC/USD	0.038	-0.045	-0.990
EUR/USD	0.026	-0.057	-0.975
EUR/JPY	0.011	-0.067	-0.943
EUR/GBP	0.001	-0.059	-0.953

We attribute the remained fluctuations to the algorithm of the price dissection into a collection of alternating trends. We also claim that the choice of the number of bins in the experiment affects the shape of the autocorrelation function in theta time. The dissection procedure has to be initialised only once and then performs unsupervised. The evolution of the price curve dictates the sequence of intrinsic events. This fact leads to the certain dilemma: to which bin of a week the intrinsic event should be assigned? A couple of prices, at which two subsequent directional changes become registered, could belong to a different bins. Let us say these are the intervals Δt_{n-1} and Δt_n. This means that the beginning of the price move that eventually reached the level of the latest intrinsic event had started within Δt_{n-1}. The end of this price trajectory finishes within the interval Δt_n. The crucial point is at what part of the Δt this beginning and end are located. In the extreme case, the whole price trajectory before the directional change could be fully placed inside of the interval Δt_{n-1}. The latest tick that eventually triggered the new directional-change event can be at the very beginning of Δt_n. Should such an event be assigned to the bin Δt_{n-1} or to Δt_n? The answer to this question is particularly important considering the effect shown in figure 6. The inconsistency affects the seasonality pattern by not only changing its average amplitude but also by providing slightly different regions of low and high instantaneous volatility (see, for example, the curves for $\delta = 0.01\%$ and $\delta = 0.10\%$).

A better way of associating locations of intrinsic events with bins of a week is another question related to the transition from the physical to intrinsic time and vice versa. This topic should be discussed in more details in further research works. Until then, a potential strategy to resolve the localisation problem is the use of smaller thresholds and bigger time intervals.

5. Concluding remarks

Analytical solutions presented in this paper translate the language of traditional risk-management tools based on drawdowns and drawups into the language of the directional-change intrinsic time. This connection makes it possible to interpret the evolution of a price curve as a sequence of alternating trends of the given scale. The observed number of directional changes for a period of time has been connected to the properties of the studied time series characterised by the instantaneous volatility σ_{DC} and the trend μ. The choice of directional-change thresholds δ_{up} and δ_{down} used to dissect the historical price curve is arbitrary but affects the results of the experiment. Equations 12, 6, and 7 have been confirmed by a Monte Carlo simulation which demonstrates the robustness and accuracy of the obtained analytical expressions.

We extended the work of Dacorogna *et al.* (1993) through discovering the weekly seasonality pattern of the instantaneous volatility of Bitcoin prices as well as of three FX exchange rates. The connection of the number of directional-change intrinsic events to the instantaneous volatility has been employed to perform the computation. Similar patterns of the realised and the instantaneous volatilities were obtained. Several noticeable differences between the results demonstrated in the work Dacorogna *et al.* (1993) and the ones presented in the current paper have been highlighted. First, the novel method significantly simplifies the construction of the instantaneous volatility seasonality pattern using tick-by-tick prices. Second, the autocorrelation function of the number of directional changes computed in physical time stays positive for a notably long period of time. Third, the beginning of the autocorrelation function computed in Θ-time can be approximated by the exponential function when the rest of it declines linearly.

The number of directional changes, directly connected to the instantaneous volatility of the given discrete time series, reveals the differences between scales of measures utilised to study the evolution of the price curve. The proposed framework represents a novel conceptual paradigm where measures are independent of the flow of physical time. We argue that this characteristic has significant advantages over 'natural' volatility estimators and has to be considered as the primary tool in the set of classical risk management techniques.

The insights provided within this paper underline the relevance of the proposed directional-change framework as a valuable alternative to the traditional time-series analysis tools. The directional-change intrinsic time has the remarkable ability to efficiently deal with fat-tailed distribution of returns and is more efficient in capturing periods of changing price activity. Results of the provided research extend the set of risk management tools constructed to evaluate the statistical properties of traditional and emerging financial markets.

References

Aı, Y., Kimmel, R. *et al.*, Maximum likelihood estimation of stochastic volatility models. *Journal of financial economics*, 2007, **83**, 413–452.

Andersen, T.G. and Lund, J., Estimating continuous-time stochastic volatility models of the short-term interest rate. *Journal of econometrics*, 1997, **77**, 343–377.

Asmussen, S. and Albrecher, H., *Ruin probabilities*, 2010, World Scientific Publishing Co Pte Ltd.

Barndorff-Nielsen, O.E. and Shephard, N., Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2002, **64**, 253–280.

Bessembinder, H., Bid-ask spreads in the interbank foreign exchange markets. *Journal of Financial economics*, 1994, **35**, 317–348.

BIS, Triennial central bank survey: Foreign exchange turnover in April 2016. , 2016.

Black, F. and Scholes, M., The pricing of options and corporate liabilities. *Journal of political economy*, 1973, **81**, 637–654.

Bollerslev, T. and Melvin, M., Bidask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics*, 1994, **36**, 355–372.

Carr, P., Zhang, H. and Hadjiliadis, O., Maximum drawdown insurance. *International Journal of Theoretical and Applied Finance*, 2011, **14**, 1195–1230.

Chekhlov, A., Uryasev, S. and Zabarankin, M., Drawdown measure in portfolio optimization. *International Journal of Theoretical and Applied Finance*, 2005, **8**, 13–58.

Cho, D.C. and Frees, E.W., Estimating the volatility of discrete stock prices. *The Journal of Finance*, 1988, **43**, 451–466.

Cont, R., Empirical properties of asset returns: stylized facts and statistical issues. , 2001.

Cotter, J., Tail behaviour of the euro. *Applied Economics*, 2005, **37**, 827–840.

Dacorogna, M.M., Müller, U.A., Nagler, R.J., Olsen, R.B. and Pictet, O.V., A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. *Journal of International Money and Finance*, 1993, **12**, 413–438.

Dacorogna, M.M., Müller, U.A., Pictet, O.V. and De Vries, C.G., Extremal forex returns in extremely large data sets. *Extremes*, 2001, **4**, 105–127.

Eross, A., Mcgroarty, F., Urquhart, A. and Wolfe, S., The Intraday Dynamics of Bitcoin. *SSRN Elecronic Journal*, 2017, pp. 1–24.

French, K.R. and Roll, R., Stock return variances: The arrival of information and the reaction of traders. *Journal of financial economics*, 1986, **17**, 5–26.

Gennotte, G. and Leland, H., Market liquidity, hedging, and crashes. *The American Economic Review*, 1990, pp. 999–1021.

Glattfelder, J., Dupuis, A. and Olsen, R., Patterns in high-frequency FX data: discovery of 12 empirical scaling laws. *Quantitative Finance*, 2011, **11**, 599–614.

Golub, A., Chliamovitch, G., Dupuis, A. and Chopard, B., Multi-scale representation of high frequency market liquidity. *Available at SSRN 2393428*, 2014.

Golub, A., Glattfelder, J. and Olsen, R.B., The alpha engine: designing an automated trading algorithm. , 2017.

Gottlieb, G. and Kalay, A., Implications of the discreteness of observed stock prices. *The Journal of Finance*, 1985, **40**, 135–153.

Grossman, S.J. and Zhou, Z., Optimal investment strategies for controlling drawdowns. *Mathematical finance*, 1993, **3**, 241–276.

Guillaume, D.M., Dacorogna, M.M., Davé, R.R., Müller, U.A., Olsen, R.B. and Pictet, O.V., From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and stochastics*, 1997, **1**, 95–129.

Hadjiliadis, O. and Večeř, J., Drawdowns preceding rallies in the Brownian motion model. *Quantitative Finance*, 2006, **6**, 403–409.

Haferkorn, M. and Diaz, J.M.Q., Seasonality and interconnectivity within cryptocurrencies-an analysis on the basis of Bitcoin, Litecoin and Namecoin. In *Proceedings of the International Workshop on Enterprise Applications and Services in the Finance Industry*, pp. 106–120, 2014.

Harris, L., *Trading and exchanges: Market microstructure for practitioners*, 2003, Oxford University Press, USA.

Hartmann, P., Trading volumes and transaction costs in the foreign exchange market: evidence from daily dollar–yen spot data. *Journal of Banking & Finance*, 1999, **23**, 801–824.

JE de Vries, J.E. and Aalborg, H.A., What can explain the price, volatility and traded volume of Bitcoin?. Master's thesis, University of Stavanger, Norway, 2017.

Jondeau, E. and Rockinger, M., Testing for differences in the tails of stock-market returns. *Journal of Empirical Finance*, 2003, **10**, 559–581.

Kelly Jr, J.L., A new interpretation of information rate. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pp. 25–34, 2011, World Scientific.

Landriault, D., Li, B. and Zhang, H., On the frequency of drawdowns for brownian motion processes. *Journal of Applied Probability*, 2015, **52**, 191–208.

Liu, R., Shao, Z., Wei, G. and Wang, W., GARCH Model With Fat-Tailed Distributions and Bitcoin Exchange Rate Returns. *Journal of Accounting, Business and Finance Research*, 2017, **1**, 71–75.

Lundberg, F., Uber die Theorie der Ruck-versicherung. In *Proceedings of the Trans VI International Congress Actuaries*, pp. 877–955, 1909.

Mancino, M.E., Recchioni, M.C. and Sanfelici, S., Estimation of Instantaneous Volatility. In *Fourier-Malliavin Volatility Estimation*, pp. 31–47, 2017, Springer.

Menkhoff, L., Sarno, L., Schmeling, M. and Schrimpf, A., Carry trades and global foreign exchange volatility.

The Journal of Finance, 2012, **67**, 681–718.

Mijatović, A. and Pistorius, M.R., On the drawdown of completely asymmetric Lévy processes. *Stochastic Processes and their Applications*, 2012, **122**, 3812–3836.

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

Nakamoto, S., Bitcoin: A peer-to-peer electronic cash system. , 2008.

Osborne, M.F., Brownian motion in the stock market. *Operations research*, 1959, **7**, 145–173.

Pospisil, L., Vecer, J. and Hadjiliadis, O., Formulas for stopped diffusion processes with stopping times based on drawdowns and drawups. *Stochastic Processes and their Applications*, 2009, **119**, 2563–2578.

Rachev, S.T., Menn, C. and Fabozzi, F.J., *Fat-tailed and skewed asset return distributions: implications for risk management, portfolio selection, and option pricing*, Vol. 139, , 2005, John Wiley & Sons.

Rolski, T., Schmidli, H., Schmidt, V. and Teugels, J.L., *Stochastic processes for insurance and finance*, Vol. 505, , 2009, John Wiley & Sons.

Sapuric, S. and Kokkinaki, A., Bitcoin Is Volatile! Isnt that Right?. In *Proceedings of the International Conference on Business Information Systems*, pp. 255–265, 2014.

Schuhmacher, F. and Eling, M., Sufficient conditions for expected utility to imply drawdown-based performance rankings. *Journal of Banking & Finance*, 2011, **35**, 2311–2318.

Taylor, H.M., A stopped Brownian motion formula. *The Annals of Probability*, 1975, pp. 234–246.

Zhang, H., Occupation times, drawdowns, and drawups for one-dimensional regular diffusions. *Advances in Applied Probability*, 2015, **47**, 210–230.

Zhang, H. and Hadjiliadis, O., Drawdowns and the speed of market crash. *Methodology and Computing in Applied Probability*, 2012, **14**, 739–752.

Appendix A: Drawdowns and drawups

Probabilities of drawdowns and drawups were extensively studied by Carr *et al.* (2011) to propose a new insurance technique against unexpected price moves and a novel way of hedging liabilities associated with these risks. Zhang and Hadjiliadis (2012) employed drawdowns as an estimate of the stock default risk and also provided a risk-management mechanism affecting the investor's optimal cancellation timing. In Schuhmacher and Eling (2011) drawdowns are considered as one of 14 reward-to-risk ratios alternative to the widely known performance measure such as the Sharpe ratio. Properties of drawdowns can be applied as an estimate of the portfolio optimisation problem (Grossman and Zhou 1993, Chekhlov *et al.* 2005). The latter can be personalised to match traders' or investors' expectations and the tolerance to the size and the length of the market disruption.

Drawdowns D_t and drawups U_t, also called rallies (Hadjiliadis and Večeř 2006), registered by the moment of time t depend on the running price maxima $\overline{P_t}$ and the running price minima $\underline{P_t}$ (Zhang and Hadjiliadis 2012, Mijatović and Pistorius 2012, Landriault *et al.* 2015). These reference points hinge on the set of historical prices P_s and are mathematically defined in the following way:

$$\overline{P_t} = \sup\{P_s : 0 \leq s \leq t\} \quad \text{and} \quad \underline{P_t} = \inf\{P_s : 0 \leq s \leq t\}, \tag{A1}$$

where $t \geq 0$ and the interval $[0, t]$ is fixed. Drawdowns and drawups are the differences between the final price of the given time interval P_t and the registered local maxima and minima:

$$D_t = \overline{P_t} - P_t \quad \text{and} \quad U_t = P_t - \underline{P_t}. \tag{A2}$$

Once a price curve experiences a drawdown D_t of the size a, the waiting time τ_a^D is registered. Similarly, for a drawup the waiting time associated with the size a is τ_a^U. In details:

$$\tau_a^D = \inf\{D_t \geq a : t \geq 0\} \quad \text{and} \quad \tau_a^U = \inf\{U_t \geq a : t \geq 0\}. \tag{A3}$$

Appendix B: Intrinsic time

All relevant to the performance of the financial system events such as political decisions, natural disasters, or economic reports rarely happen synchronously or are equally spaced in time. A sequence of them has a non-homogeneous nature and is not characterized by any vital autocorrelation function. Ultimately, the change of days and nights, as well as seasons, is dictated by the natural structure of the physical world which is barely connected with the flow of financial activity. Human minds with the whole diversity of peculiar and indescribable characteristics are primal engines of all market's evolutionary shifts. The global market, where the majority of transactions happen online and where traders, dealers, and market makers are distributed all around the world, is completely blind and deaf to the periodicity of days and nights as well as to the climate factors of any standalone region of the Earth.

Guillaume *et al.* (1997) provided the concept of alternating directional changes capable to connect the continuous flow of physical time with the evolution of price returns. In this event-based space, only a sequence of prices at which directional changes of the given size become observed and corresponding local extremes describe changes of the system's states. Thus, the set of intrinsic events is decoupled from the flow of physical time. An example of a price curve dissected by the directional-change algorithm is provided in figure B1. The density of intrinsic events depends only on the evolution of the price curve. The latter is pronounced in different number of events registered within intervals of equal length. Thus, only the end of the section 1 is located in the period T_1 while T_2 contains ends of three sections, $2, 3, 4$, and T_3 holds ends of two sections, 5 and 6. This property of directional-change intrinsic time allows to efficiently capture the most relevant to risk management information: tipping points of trend changes. At the same time, it

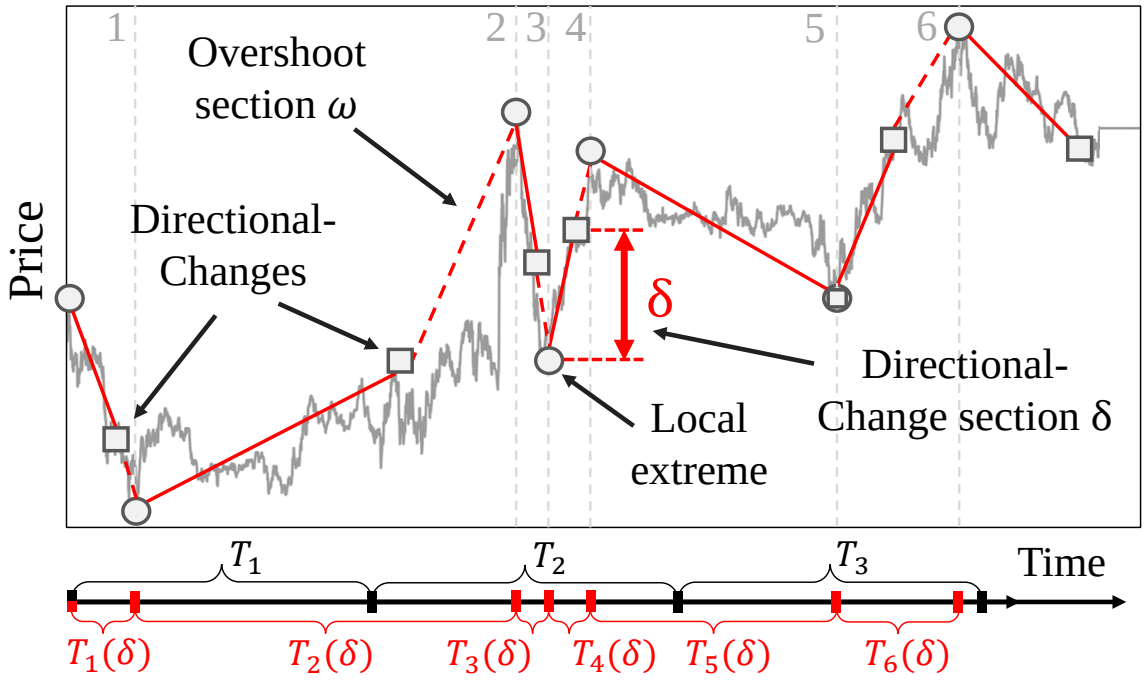

Figure B1. A price curve (grey curve) dissected into a set of directional-changes (grey squares) using a directional-change threshold δ. Grey circles mark local extremes between two consecutive directional changes. Vertical distance between each directional-change and preceding extreme price is bigger or equal to the size of the threshold δ. Vertical dashed lines indicate the end of each trend section (identified only after the next event becomes observed). The line below contains equal time intervals T_1, T_2, T_3 and length of each directional-change section in intrinsic time $T_1(\delta), \ldots, T_6(\delta)$.

possesses properties of a filtering technique by ignoring price changes between directional changes. In contrast, equally spaced through periods T_1, T_2 and T_3 prices extracted from the same curve do not contain information on its extreme activity located in T_2. This disability of the traditional techniques over stochasticity of market's speed develops into too stiff volatility estimators.

Specific properties of historical prices can be described studding a collection of directional changes. Guillaume *et al.* (1997) was the first researcher to uncover a scaling law[1] relating the expected number of directional-changes $N(\delta)$ to the size of the threshold δ. Mathematically:

$$N(\delta) = \left(\frac{\delta}{C_{N,DC}} \right)^{E_{N,DC}}, \tag{B1}$$

where $C_{N,DC}$ and $E_{N,DC}$ are the scaling law coefficients. Later, Glattfelder *et al.* (2011) employed the framework to discover 12 independent scaling laws which hold across three orders of magnitude and are present in 13 currency exchange rates. The persistence of revealed scaling laws became the base elements for the tools designed to monitor market's liquidity at multiple scales (Golub *et al.* 2014). Later Golub *et al.* (2014) described a successful trading strategy exploiting a collection of tools build upon the directional-change intrinsic time and characterized by the annual Sharp ratio greater than 3.0.

[1] A basic polynomial functional relationship where a change in input results in a proportional change in output.

Appendix C: Bitcoin seasonality

Despite the relatively young age of the blockchain technology, there are already a few research works concern the statistical properties of cryptocurrencies. Sapuric and Kokkinaki (2014) analysed realised volatility of Bitcoin returns within 4-years time interval to understand what are the prime characteristics of its price activity. They confirmed Bitcoin's high volatility, but emphasized that traded volume is a remarkable fact which should be taking into account computing the realized volatility[2]. The authors compare Bitcoin with conventional financial instruments, including gold, and several national currencies and demonstrate that the calculated volatility significantly decreases when the traded volumes are included in the model.

Haferkorn and Diaz (2014) studied seasonality patterns of the number of payments performed in three cryptocurrencies: Bitcoin (classified as a worldwide payment system), Litecoin (open source software project), and Namecoin (decentralised name system). Their research confirmed that in contrast to the traditional equity and FX markets, the monthly or yearly seasonality is not typical for the crypto market. The only robust weekly pattern was fund in Bitcoin prices while Litecoin and Namecoin had weak or no patterns at all. The authors state that there is also no significant correlation between the returns of observed exchange rates. Authors speculate that the reason for this is the fact that although these cryptocurrencies have similar core architecture, they all have been created to serve quite specific needs.

JE de Vries and Aalborg (2017) made another attempt to discover seasonality patterns of Bitcoin prices analysing daily traded volume, daily transaction volume, and Google trends (the number of searches for the word 'bitcoin'). The author also inspected the seasonality of the number of transactions performed from and to individual blockchain accounts. All of the measurements demonstrated no particular periodicity.

Eross *et al.* (2017) gave a more affirmative answer on the existence of the intraday seasonality of Bitcoin prices. The authors investigated Bitcoin returns, volume, realised volatility, and bid-ask spreads to reveal several intraday stylized facts. A significant negative correlation was found between returns and volatility, while volume and volatility have a considerably positive correlation. The authors attribute such patterns to the European and North American traders as well as to the lack of market makers in the whole crypto space.

Appendix D: Monte Carlo waiting times

Appendix E: 'Traditional' volatility seasonality

[2]According to the Bank for International Settlements the daily average FX trading volume in April 2016 was $5.1 trillion (BIS 2016) when the highest registered volume in the crypto market is only $45.8 billion (https://coinmarketcap.com/charts/).

Table D1. Waiting times and number of directional changes in a Monte Carlo simulation. μ and σ are parameters of the Brownian motion used for the test. 10^9 ticks in the simulated time series. N_{DC}^{MC}, $\langle T_{up}^{MC} \rangle$, and $\langle T_{down}^{MC} \rangle$ are the number of directional changes and the average waiting times registered in the Monte Carlo simulation. $\mathbb{E}[N_{DC}]$, $\mathbb{E}[T_{up}]$, and $\mathbb{E}[T_{down}]$ are theoretical values dictated by equations 12, 6, and 7 correspondingly. Values $\sigma^{-}_{T_{up}^{MC}}$ and $\sigma^{-}_{T_{down}^{MC}}$ are standard deviations of empirical and theoretical waiting times.

μ, %	σ, %	$N_{DC}^{MC}/\mathbb{E}[N_{DC}]$	$\langle T_{up}^{MC} \rangle/\mathbb{E}[T_{up}]$	$\sigma^{-}_{T_{up}^{MC}}$	$\langle T_{down}^{MC} \rangle/\mathbb{E}[T_{down}]$	$\sigma^{-}_{T_{down}^{MC}}$
1	10	1.028	0.968	2.54e-05	1.019	2.53e-06
	20	1.009	0.989	2.78e-06	1.012	3.32e-07
	30	1.001	0.995	8.79e-07	1.033	9.58e-08
6	10	1.021	0.971	2.29e-05	1.043	2.59e-06
	20	1.005	0.993	2.94e-06	1.019	3.29e-07
	30	0.987	1.011	8.84e-07	1.034	9.98e-08
11	10	1.029	0.968	2.20e-05	1.011	2.78e-06
	20	0.994	1.006	2.72e-06	0.997	3.30e-07
	30	0.986	1.014	8.82e-07	1.017	1.02e-07

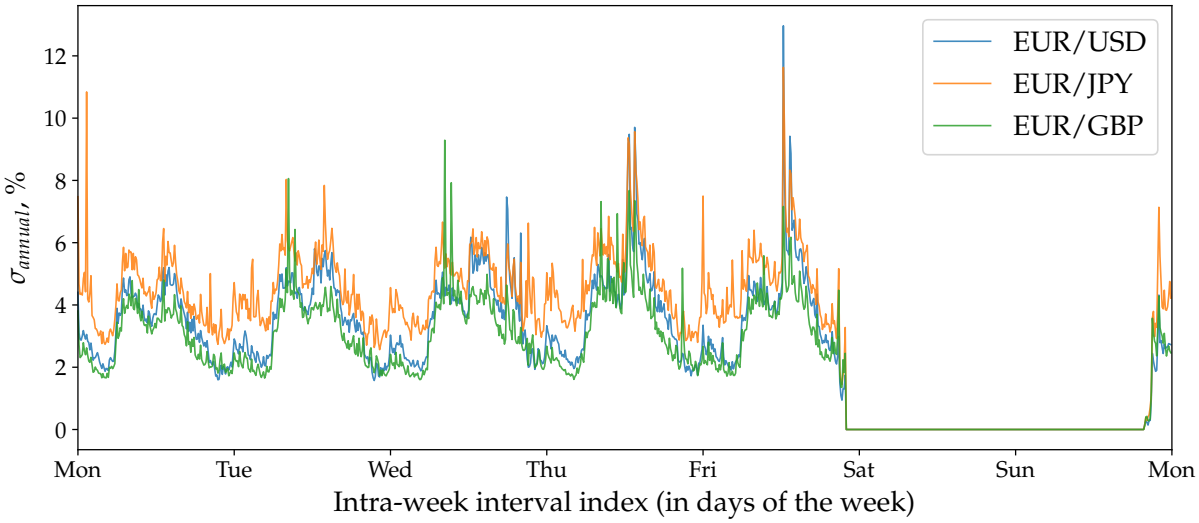

Figure E1. Realized volatility seasonality patterns of three FX exchange rates computed using the traditional approach (equation 16). 1 minute intervals have been used to calculate returns. The size of each bin is 10 minutes, 1008 bins in total.

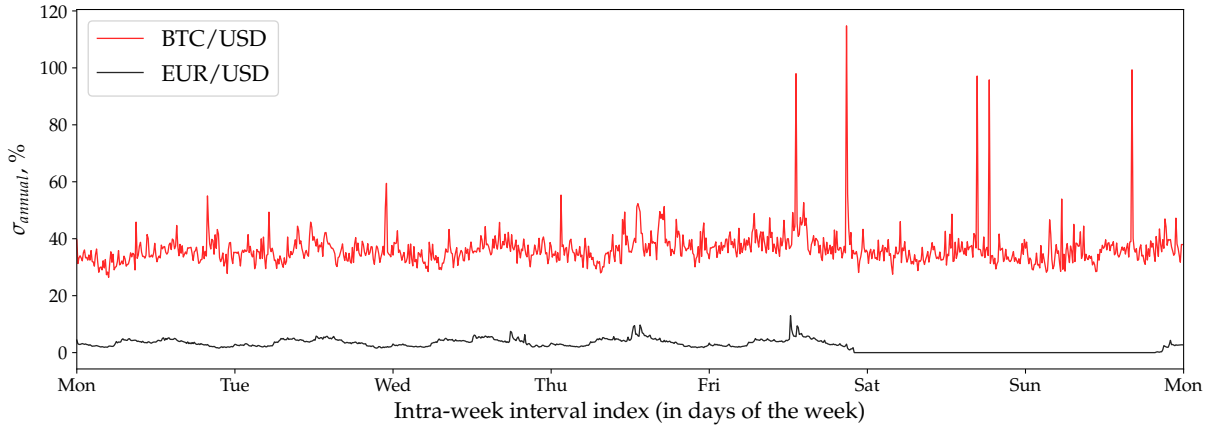

Figure E2. Realized volatility seasonality patterns of BTC/USD and EUR/USD exchange rates computed using the traditional approach (equation 16). 1 minute intervals have been used to calculate returns. The size of each bin is 10 minutes, 1008 bins in total.