



Essex Business School

Periodicities of FX Markets in Intrinsic Time

Dr Nick Constantinou
University of Essex

Iacopo Giampaoli
University of Essex

Wing Long Ng
University of Essex

Address for correspondence

Wing Long Ng
CCFEA
University of Essex
Colchester
Essex CO4 3SQ, UK

E-mail: wlng@essex.ac.uk

April, 2010

Periodicities of FX Markets in Intrinsic Time

Wing Lon Ng* Iacopo Giampaoli † Nick Constantinou ‡

April 15, 2010

***Corresponding author.** E-mail: wlng@essex.ac.uk. Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK.

†Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK.

‡Essex Business School, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK.

The authors are grateful to the Editor and an anonymous referee for their helpful comments and suggestions that led to an improvement of the paper. They are also grateful to Richard Olsen for his valuable helpful discussion and to OFT for providing the FX market data.

Periodicities of FX Markets in Intrinsic Time

Abstract

This paper utilises advanced methods from Fourier Analysis in order to describe financial ultra-high frequent transaction data. The Lomb-Scargle Fourier Transform is used to take into account the irregularity in spacing in the time-domain. It provides a natural framework for the power spectra of different inhomogeneous time series processes to be easily and quickly estimated, without significant computational effort, in contrast to the common econometric approaches in the finance literature. An event-based approach (intrinsic time), which by its own nature is inhomogeneous in time, is employed using different event thresholds to filter the foreign exchange tick-data leading to a power-law relationship. The calculated spectral density demonstrates that the price process in intrinsic time contains different periodic components, especially in the medium-long term, implying the existence of new stylised facts of ultra-high frequency data in the frequency domain.

Keywords: Ultra-high frequency transaction data, foreign exchange, irregularly spaced data, Lomb-Scargle Fourier Transforms, spectral density.

JEL Classifications: C22, C46, C63, G17.

1 Introduction

In financial markets, trading activity tends typically to vary depending on the time of day, shrinking during lunch time and over weekends, and increasing prior to major scheduled news announcements (see, for example, [Chordia, Roll, and Subrahmanyam, 2001](#)). These observations are usually measured in *physical time*. Another approach that has been recently put forward and which has led to a rich array of scaling behaviour in FX market is a time scale that is defined via events, i.e. the so called *intrinsic time* (see also [Glattfelder, Dupuis, and Olsen, 2009](#); [Bisig, Dupuis, Impagliazzo, and Olsen, 2009](#)).

In particular, the presence of different patterns of trading activity in financial markets makes the flow of physical time discontinuous, as the number of transactions tends to increase or decrease at different time intervals (for example, over weekends). This empirical observation, which led to the definition of “intrinsic trading time”, coined by [Mandelbrot and Taylor \(1967\)](#), suggests that the concept of physical time may not be the fundamental time scale and research should additionally focus on an event-based time scale. Thus, the analysis of price movements investigated in this paper, focuses on an intrinsic time scale defined by “directional-change” events: price movements exceeding a given threshold which are independent of the notion of physical time. Indeed, the recent work of [Glattfelder, Dupuis, and Olsen \(2009\)](#) has demonstrated a rich landscape of scaling laws in intrinsic time for “ultra-high-frequency” (UHF) FX data, including a scaling law for the directional-change events considered in this paper.

In the intrinsic time paradigm, the UHF data are filtered by analysing those points where there has been a directional-change (the *event*), consisting of price

movements away from a local maximum (or minimum), which exceed a given threshold (for a mathematical definition of a directional change duration, see [Hautsch, 2004](#)). The newly created time series naturally includes fewer observations, but still contains significant information about the price evolution via its associated scaling law and about the temporal structures through its directional-change duration, and is irregular-spaced in time. Indeed, time series in intrinsic time are fundamentally unevenly-spaced in time (“inhomogeneous”) and there is no justification in artificially making them equally-spaced (“homogenous”) (see also [Dacorogna, Gençay, Müller, Olsen, and Pictet, 2001](#)).

The concept of directional change also plays an important role in technical analysis ([Hautsch, 2004](#)), as turning points of fundamental price movements are used by intraday traders to identify optimal times to buy or sell. For example, point-and-figure charting techniques, reportedly in use from the late 19th century, track price changes of fixed thresholds and significant price reversals, and are often employed to identify price trends, support and resistance levels in intraday trading data (among the oldest written references are [Wyckoff, 1910](#) and [deVilliers, 1933](#)). Additionally, the so called “directional change frequency”, which estimates the average number of directional price changes of a given threshold over a data sample, can be interpreted as an alternative measure of risk ([Guillaume, Dacorogna, Davé, Müller, Olsen, and Pictet, 1997](#)).

The main goal of this study is to detect potential periodic patterns of UHF FX market data. For example, it is well-known that intraday data have a consistent diurnal pattern of trading activities over the course of a trading day due to institutional characteristics of organised financial markets, such as opening and closing hours or intraday auctions. In particular, FX markets have stronger seasonal pat-

terns as they are open 24/7 and constantly influenced by re-occurring time-zone effects.

In particular, this paper employs advanced modelling techniques from Fourier analysis, which provide a natural framework to analyse inhomogeneous time series in the frequency domain and also reduce the computational time needed to process a large amount of transaction data. Fast Fourier transform (FFT) algorithms are commonly used to analyse homogeneous time series in the frequency domain, but in their standard form, they require either regular resampling of the unevenly-spaced data, or interpolating them onto an equally-spaced grid. In fact, any such transformation via regular resampling of unevenly-spaced data or interpolation to an evenly-spaced grid (in order to calculate the SDF with the simple FFT), has been demonstrated to introduce artefacts in the data in the frequency domain (and hence time domain), leading to loss of information and the use of spurious information (Giampaoli, Ng, and Constantinou, 2009). The Lomb-Scargle Fourier transform (LSFT), a generalisation of the discrete Fourier transform, is especially designed for unevenly-spaced data and represents the natural tool to study UHF data in the frequency domain as it allows the stochastic behaviour of every single process to be determined without any loss of information. The LSFT, in contrast to existing autoregressive conditional models (e.g., ACD or ACI models) in the literature, has also the advantage of greatly reducing the computational effort required when analysing UHF transaction data sets and of avoiding complex model specifications or obligatory deseasonalisation.

The paper is structured as follows. Section 2 outlines methods and models used for the analysis. Section 2.1 delineates the intrinsic-time framework, while in Section 2.2 the Lomb-Scargle Fourier transform and the estimation of the spectral

density are characterised. Data and empirical results, with regard to their economic implications, are presented in Section 3. Section 4 concludes.

2 Methodology

In this paper, an analysis is made in the frequency domain of an intrinsic-time process defined by a directional-change event. In the following section, the adopted event-based approach and the resulting scaling law are described in detail. The Lomb-Scargle Fourier Transform for the spectral analysis of limit order-book data is then introduced and defined.

2.1 Intrinsic Time and Physical Time

Let x denote the price and Δx the price change. In the event-driven setting outlined above, the event is a directional change and its accompanied overshoot. More precisely, the absolute price change (“total price movement”) Δx_{tm} between two local extremal values (minimum and maximum price) is decomposed into two sections: a fixed directional change threshold Δx_{dc} and an overshoot Δx_{os} , which represents the price movement beyond the fixed threshold. Such fixed event threshold defines the minimum price move between two consecutive local extrema and fully characterises the directional changes, which are computed iteratively from the last high or low, depending on whether the directional change is assumed to be downward or upward. At each iteration, the high is updated to the maximum between the current price and the last high, or the low to the minimum between the current price and the last low. At each occurrence of a directional-change event, the direction of the price move alternates, and the overshoot Δx_{os} associated with the previous directional change

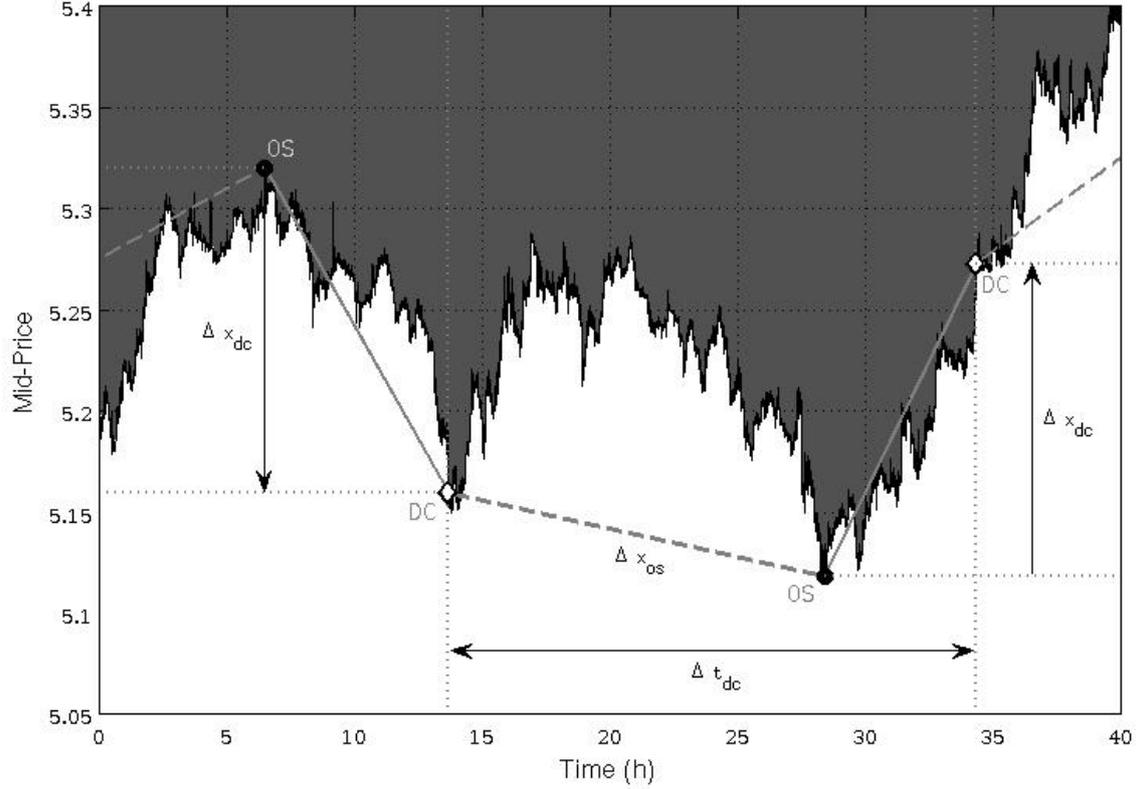


Figure 1: Dissection of mid-price curve into directional-change and overshoot sections. The graph shows a 40-hour mid-price sample and corresponding directional-change events defined by a threshold $\Delta x_{dc} = 0.3\%$ (currency pair: AUD-HKD). The directional-change events (diamonds) act as cut-off points, decomposing an absolute price change between a minimum and a maximum price (bullets) into directional-change Δx_{dc} (solid lines) and overshoot Δx_{os} (dashed lines) components. Intrinsic time ticks only at directional-change events, and the elapsing time between two contiguous events, i.e. the directional change duration, is denoted by Δt_{dc} .

is determined *ex-post* as the difference between the last high or low and the price corresponding to that directional change. A price sample and its decomposition into directional-change and overshoot components are shown in Figure 1.

Let $\langle \cdot \rangle$ denotes the average operator. [Glattfelder, Dupuis, and Olsen \(2009\)](#) have shown that the mean absolute price movement can be decomposed into a directional change and an overshoot, i.e. $\langle |\Delta x_{tm}| \rangle = \langle |\Delta x_{dc}| \rangle + \langle |\Delta x_{os}| \rangle$. The points in which the threshold is triggered, either upwards or downwards, are independent of the

notion of physical time and form a series of directional-change events to be analysed using the LSFT. This event is chosen, not only because of (a) its scaling properties but also for (b) the “natural” periodic behaviour of FX markets due to their market microstructure effects (see, e.g., [Lyons \(2001\)](#) or [MacDonald \(2007\)](#)) .

Let t_{dc} denote the directional change duration, i.e. the elapsed time of a directional change of a price movement for a given threshold size. The model applied here is the empirical scaling law obtained by [Glattfelder, Dupuis, and Olsen \(2009\)](#) that describes

$$\langle \Delta t_{dc} \rangle = c(\Delta x_{dc})^k \quad (1)$$

where c is a constant and k is the scaling exponent. This scaling law investigates the relationship between the average duration of directional changes (which is a random variable) and the size of a given price move threshold (which is pre-specified). By taking the logarithm of both sides of eq. (1), then the power law relationship is cast into the linear equation

$$\log(\langle \Delta t_{dc} \rangle) = \log(c) + k \cdot \log(\Delta x_{dc}) \quad , \quad (2)$$

characterised by the slope k and the intercept $\log(c)$. The scaling exponent k measures the proportional change of the average directional-change duration due to an increment of the price threshold. Different directional-change threshold sizes $\Delta x_{dc} = \{10\text{bp}, 15\text{bp}, 20\text{bp}, 25\text{bp}, 30\text{bp}, 40\text{bp}, 50\text{bp}, 75\text{bp}, 100\text{bp}, 125\text{bp}, 150\text{bp}, 175\text{bp}, 200\text{bp}\}$ are chosen to generate the respective average directional-change durations for the individual sampling window. Standard linear OLS regression is used

to estimate the scaling law parameters in eq. (2).

As discussed, the cut-off points between each directional change and the corresponding overshoot form a new event-based and irregularly-spaced time series in intrinsic time (see Figure 1), which serves as input for the LSFT to calculate the spectral density function (SDF). The LSFT framework is now described theoretically.

2.2 Spectral Analysis of Tick Data

In the last decade, the literature on financial econometrics has witnessed a growing interest in the analysis of tick-by-tick transaction data provided by electronic trading platforms. Since these UHF data are observed in real-time, yielding the highest possible sampling limit, they are characterised by the irregularity of time intervals between two consecutive events. Duration, defined as elapsing time between two successive orders, is in fact a crucial variable as it reflects the intensity of trading activity and thus different levels of the asset's liquidity. Therefore, the econometric analysis of unevenly-spaced transaction data has focused on modelling the durations in order to avoid any loss of important information stored in the temporal structure of the transaction process. Treating these special time series as point processes, [Engle and Russell \(1998\)](#) have introduced the Autoregressive Conditional Duration (ACD) model, which describes a dynamic duration process with a conditional expectation written as a linear function of past durations (for a survey, see [Pacurar, 2008](#) or [Bauwens and Hautsch, 2009](#)). In addition, following [Cox and Isham \(1980\)](#), alternative approaches to deal with ultra-high frequency data as point processes have been developed such as count models (see [Heinen and Rengifo, 2007](#))

or intensity models (see [Hall and Hautsch, 2006](#), [Bauwens and Hautsch, 2009](#)).

Although these three types of models have been improved by many authors and have shown good performance in numerous previous studies, they still have drawbacks and limits. Duration models, for example, can not be extended to multidimensional settings due to the asynchronisation problem of multivariate point processes ([Hall and Hautsch, 2006](#)). In contrast, other approaches can overcome this problem, but they either lose information because of the aggregation over discrete time intervals (count models) or induce computational complexity by mimicking continuous time models (intensity models). Since point processes naturally focus on the time of the actual observation (the “points”), the “marks” (i.e. the observation itself) usually only serve as regressors explaining the duration or the intensity of the process. Alternatively, the relationship between different covariates can be investigated by decomposition methods (see [Engle, 2000](#) or [Rydberg and Shephard, 2003](#)). However, the stochastic process of the key variables are never modelled as single self-contained processes, but as subordinators.

In order to address the problems mentioned above, this paper suggests using the LSFT as the natural setting to analyse unequally-spaced time series in the frequency domain. This approach has the undoubted advantage of reducing the computational load required by standard models to process large amounts of data and of avoiding any data manipulation that would alter and impair the information contained in the UHF data ([Giampaoli, Ng, and Constantinou, 2009](#)).

Spectral analysis decomposes a time series into its periodic frequencies in order to detect and analyse its cyclical behaviour. The application of spectral techniques in periodic economic processes has a long history, and a significant effort is employed in the estimation of the spectral density function (SDF) (see [Priestley, 1981](#); a survey

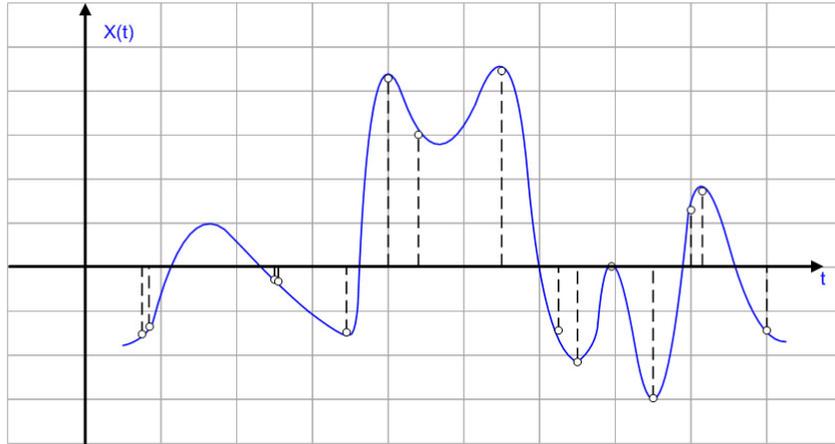


Figure 2: Fitting the FFT to unevenly-spaced UHF time series data.

of spectral analysis of economic time series can be found in [Granger and Engle, 1984](#) and [Iacobucci, 2005](#)). The SDF represents the analogue of the autocorrelation function in the time domain and encapsulates the frequency properties of the time series determining how the variation in a time series is built-up by components at different frequencies.

For standard periodic time series a FFT algorithm is generally employed to determine their spectral properties ([Priestley, 1981](#)). As tick-by-tick data and the resultant filtered time series in intrinsic time are unevenly-spaced, the traditional FFT can not be applied without artificially tampering with the raw data. Attempts to transform the irregularly-spaced raw data into regularly-spaced data (e.g., [Dacorogna, Gençay, Müller, Olsen, and Pictet, 2001](#)), by regular resampling or using interpolation, prior to applying the FFT to calculate the SDF, has recently been demonstrated to cause (a) loss of information and (b) generation of spurious data. It has also been shown that these limitations can be overcome using the LSFT (see [Giampaoli, Ng, and Constantinou, 2009](#)). [Lomb \(1976\)](#) first introduced this statisti-

cal method in astrophysics and fitted sinusoidal curves to the unevenly-spaced data (for illustration, see Figure 2) by least-squares in order to determine their periodic behaviour, despite the irregular spacing in time. Scargle (1982) later developed the methodology further by deriving the standardised Lomb-Scargle (LS) periodogram that has well defined statistical properties as the work of Horne and Baliunas (1986) demonstrated. Press and Rybicki (1989) suggested an alternative algorithm for a faster and more efficient computation (for its implementation see Press, Flannery, Teukolsky, and Vetterling, 1992). The mathematical properties of the FFT and the LSFT are now briefly outlined.

A finite time series x_t with length T and N observations is considered. For common time series, the time interval between two observations is constant, i.e. $t_j - t_{j-1} = \Delta t = T^{-1}$, $j \in \{1, 2, 3, \dots, N\}$. In Fourier analysis the time series can be expressed as a sum of trigonometric functions

$$x_t = \sum_{k=-N/2}^{N/2-1} (a_k \cos(\omega_k t) + b_k \sin(\omega_k t)) = \sum_{k=-N/2}^{N/2-1} c_k e^{i\omega_k t} \quad (3)$$

where the angular frequency is $\omega_k = 2\pi k/N$ and the frequency $f_k = \omega_k/2\pi$. The normalised spectral density function (SDF) is defined as

$$SDF_{FFT}(\omega_k) = \frac{1}{N\sigma_x^2} \left| \sum_{t=1}^N x_t e^{-i\omega_k t} \right|^2 \quad (4)$$

where $\sigma_x^2 = (N-1)^{-1} \sum_{j=1}^N (x_j - \bar{x})^2$ and the coefficients $c_k = N^{-1} \sum_{t=1}^N x_t e^{-i\omega_k t}$ can be computed using the well-known FFT (see, for example, Bloomfield, 2000 and Warner, 1998). However, as mentioned above, tick data and the resultant time series, filtered in intrinsic time via the event, arrive in irregular time intervals, where

Δt is now stochastic. Hence, the simple FFT can not be employed without data manipulation. The main objective here is to find an algorithm that can compute eq. (4) fast.

Many diverse areas of science have tackled this issue using the robust framework of the LSFT (for an overview, see [Ware, 1998](#)). Under this framework, the data on the non-equally spaced grid is transformed into the frequency domain in order to obtain an unbiased estimation of the SDF. The resulting SDF is calculated for $k \in \{1, 2, 3, \dots, M\}$ frequencies, with M chosen as outlined in [Press, Flannery, Teukolsky, and Vetterling \(1992\)](#). The normalised SDF is given by

$$SDF_{LS}(\omega_k) = \frac{1}{2\sigma_x^2} \left\{ \frac{\left[\sum_{j=1}^N (x_j - \bar{x}) \cos \omega_k (t_j - \tau) \right]^2}{\sum_{j=1}^N \cos^2 \omega_k (t_j - \tau)} + \frac{\left[\sum_{j=1}^N (x_j - \bar{x}) \sin \omega_k (t_j - \tau) \right]^2}{\sum_{j=1}^N \sin^2 \omega_k (t_j - \tau)} \right\} \quad (5)$$

with $\bar{x} = N^{-1} \sum_{j=1}^N x_j$ and

$$\tau(\omega_k) = \frac{1}{2\omega_k} \arctan \left(\frac{\sum_{j=1}^N \sin(2\omega_k t_j)}{\sum_{j=1}^N \cos(2\omega_k t_j)} \right) \quad (6)$$

with $f_k = \omega_k/2\pi \in [0, 0.5]$ as the frequency (see [Press, Flannery, Teukolsky, and Vetterling, 1992](#), p. 581; a generalisation for the non-sinusoidal case can be found in [Bretthorst, 2001](#)). Due to the (re-)shifting of all N sampling times t_j with τ , the time invariance of $f(\omega_k)$ is ensured. In addition, [Scargle \(1982, Appendix C\)](#) demonstrated that this particular choice of the offset τ makes eq. (5) identical to the expression obtained by linear least-squares fitting sine waves to the data

(see also [Van Dongen, Olofsen, Van Hartevelt, and Kruyt, 1999](#)). [Scargle \(1982\)](#) has also shown that the Lomb-Scargle periodogram has an exponential probability distribution with unit mean. The probability that SDF_{LS} will be between some positive quantity z and $z + dz$ is $e^{-z} dz$, and the probability of none of them give larger values than z is $(1 - e^{-z})^M$. Therefore, we can compute the false-alarm probability of the null hypothesis, e.g., the probability that a given peak in the periodogram is not significant, by $P(> z) \equiv 1 - (1 - e^{-z})^M$ ([Press and Rybicki, 1989](#)).

3 Empirical Data and Results

The objective of this work is to highlight the advantages of the LSFT in obtaining the SDF for inhomogeneous UHF data and to demonstrate that the use of LSFT within an event-based framework reveals new periodic patterns in FX time series and provides insightful information on the price process.

The tick-by-tick data set comprises 6 currency pairs spanning 3 months, from November 1, 2008 to January 31, 2009. The following currency pairs are considered (with the number of observations enclosed in brackets): AUD-HKD (4'472'222), AUD-JPY (18'821'980), EUR-JPY (32'250'932), EUR-USD (23'057'152), HKD-JPY (6'052'923), USD-JPY (19'010'622). The varying number of ticks is mostly due to the fact that different exchange rates have different degrees of liquidity ([Glattfelder, Dupuis, and Olsen, 2009](#)). The data set includes a bid, an ask price, a timestamp, and each time series is filtered as observations with the same timestamp are averaged

Currency	AUD-HKD	AUD-JPY	EUR-JPY	EUR-USD	HKD-JPY	USD-JPY
Threshold	$N(\Delta x_{dc})$ (DC/h)					
10 bp	22425 (14.8363)	32313 (21.7158)	15557 (10.4550)	8154 (5.4831)	8564 (5.7606)	8573 (5.7614)
15 bp	10615 (7.0228)	15006 (10.0847)	7693 (5.1700)	3921 (2.6366)	3837 (2.5810)	3807 (2.5585)
20 bp	6289 (4.1608)	8947 (6.0128)	4601 (3.0921)	2300 (1.5466)	2270 (1.5269)	2241 (1.5061)
25 bp	4185 (2.7688)	5994 (4.0282)	3027 (2.0343)	1545 (1.0389)	1428 (0.9605)	1402 (0.9422)
30 bp	2962 (1.9596)	4390 (2.9503)	2157 (1.4496)	1084 (0.7289)	1023 (0.6881)	1009 (0.6781)
40 bp	1677 (1.1095)	2572 (1.7285)	1246 (0.8374)	622 (0.4183)	588 (0.3955)	578 (0.3884)
50 bp	1076 (0.7119)	1689 (1.1351)	793 (0.5329)	380 (0.2555)	379 (0.2549)	372 (0.2500)
75 bp	464 (0.3070)	754 (0.5067)	355 (0.2386)	175 (0.1177)	165 (0.1110)	162 (0.1089)
100 bp	251 (0.1661)	425 (0.2856)	188 (0.1263)	97 (0.0652)	87 (0.0585)	84 (0.0565)
125 bp	173 (0.1145)	267 (0.1794)	122 (0.0820)	57 (0.0383)	50 (0.0336)	49 (0.0329)
150 bp	121 (0.0801)	187 (0.1257)	84 (0.0565)	39 (0.0262)	30 (0.0202)	29 (0.0195)
175 bp	91 (0.0602)	133 (0.0894)	64 (0.0430)	27 (0.0182)	23 (0.0155)	22 (0.0148)
200 bp	67 (0.0443)	113 (0.0759)	54 (0.0363)	25 (0.0168)	19 (0.0128)	18 (0.0121)

Table 1: Number of directional changes (time period: 11/2008 - 01/2009). The table shows for each given threshold Δx_{dc} the total number of directional changes $N(\Delta x_{dc})$, and the number of directional changes per hour (DC/h).

out. Throughout the paper, the following definition of mid-price is considered:

$$x_t = (bid_t + ask_t) / 2. \quad (7)$$

From the raw data, a mid-price process for each of the currency pairs is obtained using eq. (7). It is the mid-price data which are filtered to intrinsic time series. The filtering follows the procedure outlined above, i.e. by decomposing the total

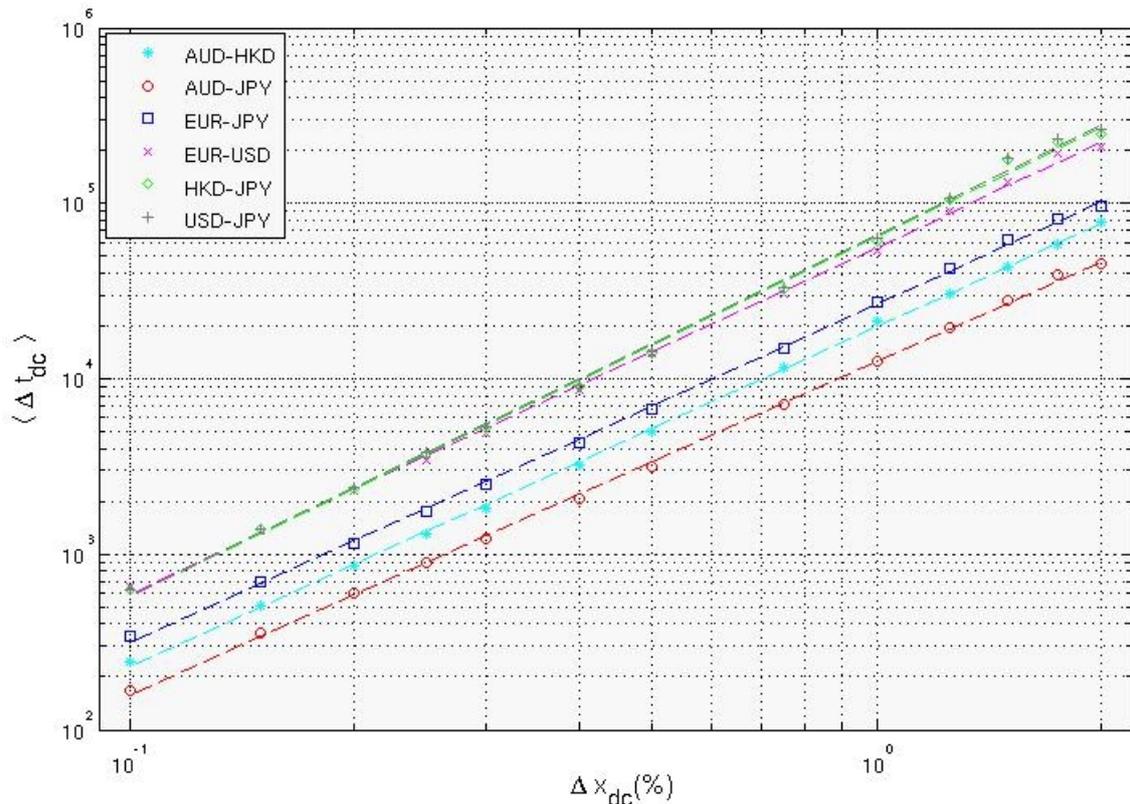


Figure 3: Empirical scaling law: estimated scaling law regression lines for the different currency pairs (time period: 11/2008 - 01/2009). The x -axis shows the threshold size as relative price change, and the y -axis the respective average time (in seconds), for a given threshold, to observe the reversion of the price move.

price movements into directional-change (of thresholds $\Delta x_{dc} = \{10\text{bp}, 15\text{bp}, 20\text{bp}, 25\text{bp}, 30\text{bp}, 40\text{bp}, 50\text{bp}, 75\text{bp}, 100\text{bp}, 125\text{bp}, 150\text{bp}, 175\text{bp}, 200\text{bp}\}$) and overshoot sections. Table 1 shows, for each threshold Δx_{dc} , the number of directional changes $N(\Delta x_{dc})$ and the “speed” of change of the mid-price process, measured in directional changes per hour (DC/h).

Figure 3 illustrates the power law regression (eq. 2) for the currency pairs considered in this paper. Table 2 lists the intercept, slope, associated R^2 and mean square error (MSE) statistics. For each of the currency pairs, we show the results

Currency	Intercept (s.e.)	Slope (s.e.)	R^2	MSE
AUD-HKD	8.1806 (0.0258)	1.9407 (0.0111)	0.9996	2.7E-4
GBM	5.9829	1.7176	0.9965	0.0021
AUD-JPY	7.8952 (0.0284)	1.8985 (0.0123)	0.9995	3.3E-4
GBM	5.6690	1.6578	0.9952	0.0027
EUR-JPY	8.2895 (0.0342)	1.9316 (0.0148)	0.9994	4.8E-4
GBM	6.1383	1.7451	0.9970	0.0019
EUR-USD	8.7042 (0.0445)	1.9768 (0.0192)	0.9990	8.2E-4
GBM	6.5774	1.8115	0.9978	0.0015
HKD-JPY	8.9131 (0.0602)	2.0520 (0.0260)	0.9982	0.0015
GBM	6.5938	1.8149	0.9978	0.0015
USD-JPY	8.9539 (0.0604)	2.0656 (0.0260)	0.9983	0.0015
GBM	6.6002	1.8158	0.9978	0.0015

Table 2: Empirical scaling law: estimated regression parameters (time period: 11/2008 - 01/2009). The scaling law relates the average time interval for directional changes of given thresholds to occur to the size of the thresholds. For each currency pair, the last row shows the regression parameters of the benchmark geometric Brownian motion (GBM).

for a geometric Brownian motion (GBM), used as a benchmark. Given an initial mid-price x_0 , the process generated is

$$x_t = x_0 \exp \left(\left(\mu - \frac{\sigma^2}{2} \right) t + \sigma W_t \right), \quad (8)$$

where W_t is a Wiener process, and μ and σ are estimated from the data. The regression coefficients for the GBM are obtained by averaging the OLS estimates over a sufficiently high number of iterations, so as to ensure good convergence properties. The slope of the curves (the key parameter in the power-law) are close to those reported in [Glattfelder, Dupuis, and Olsen \(2009\)](#), with similar associated statistics.

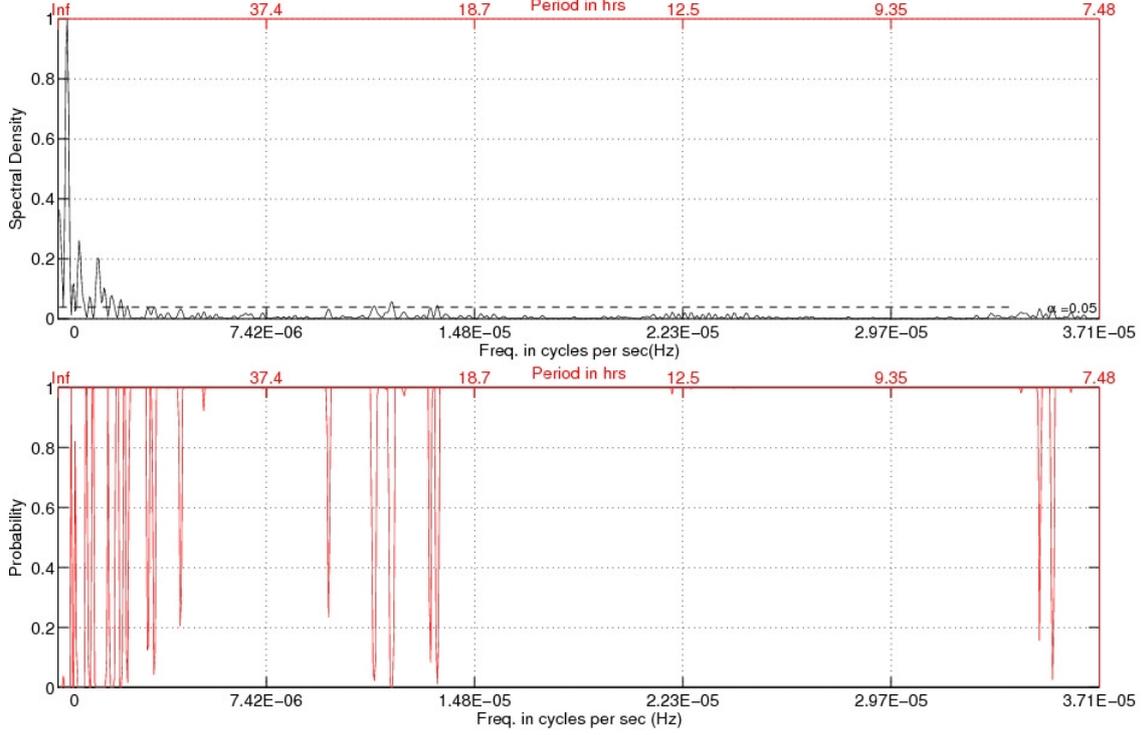


Figure 4: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 245$) of mid-price directional-change events of threshold 0.5% for the EUR-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% significance level for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis label. The graph shows that the most significant contribution to the variance of the process comes from relatively low frequencies (above 8 hours).

This similarity illustrates the remarkable robustness of the scaling law as the data used in this investigation are later than those employed by [Glattfelder, Dupuis, and Olsen \(2009\)](#). In particular, it is noted that all the currency pairs show power-law behaviour which is statistically different from the corresponding GBM at the 99% confidence level.

Following the data filtering, the dissection (cut-off) points between directional-change and overshoot sections generate a new irregularly-spaced time series of directional-change events in intrinsic time; the LSFT is then applied to calculate

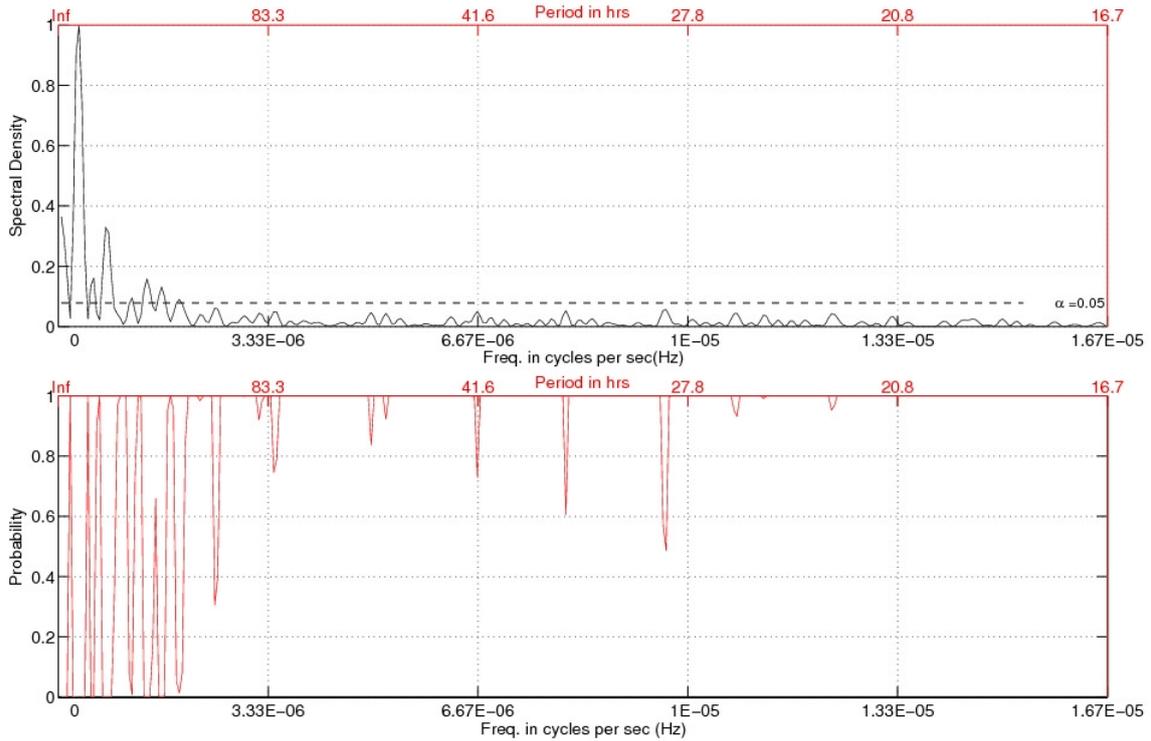


Figure 5: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 112$) of mid-price directional-change events of threshold 0.75% for the EUR-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% significance level for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis label. The graph shows that, as the threshold size increases, the SDF tends to shift towards lower frequencies (period above 100 hours).

the SDF.

Figures 4, 5, 6, show examples of the SDFs of mid-price directional-change events (in the top panels) and the corresponding false-alarm probabilities (in the bottom panels) of the currency pair EUR-JPY for 3 different thresholds (0.5%, 0.75%, 1.5%). For a better comparison, the spectral densities are normalised by the maximum of SDF_{LS} . Sample graphs for a selection of two thresholds (0.75%, 1.5%), for the remaining currency pairs and the same time period, can be found in the appendix.

Figure 4 clearly illustrates that all the significant peaks (for a significance level

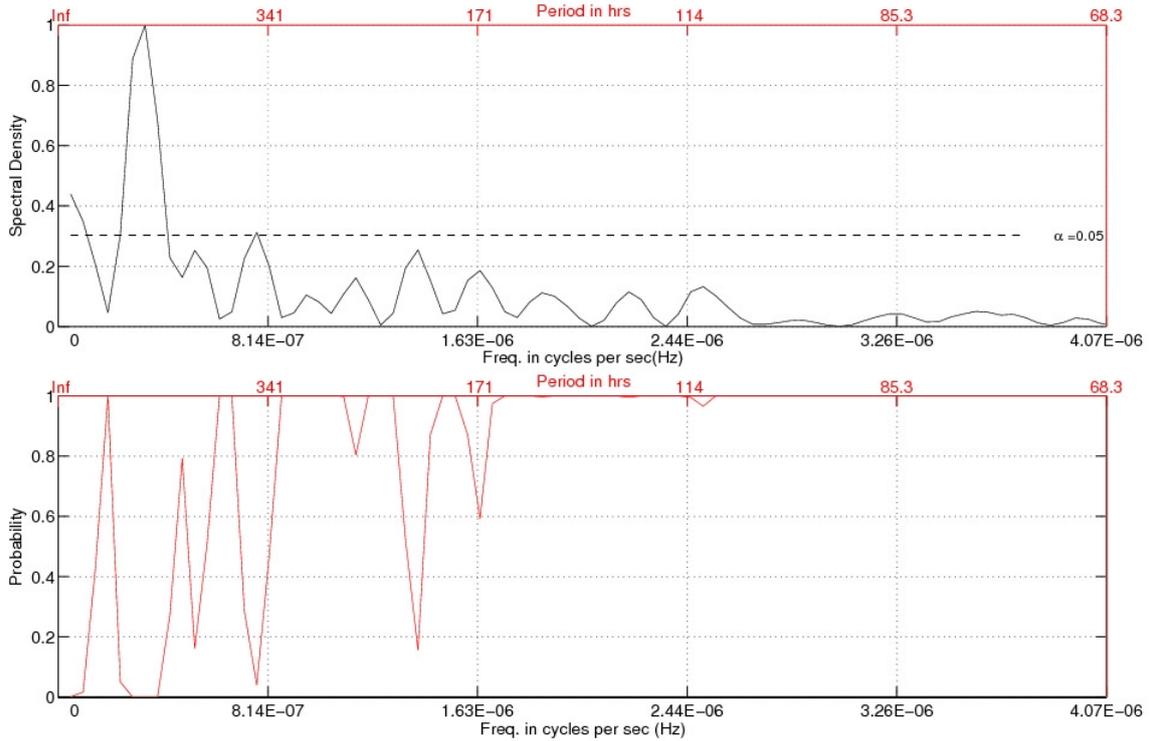


Figure 6: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 24$) of mid-price directional-change events of threshold 1.5% for the EUR-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% significance level for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis label. As the graph illustrates, the SDF continues shifting towards the left-hand side (period above 340 hours).

of 95%) are located in the left-hand side of the graph¹, i.e. that the highest contribution to the variance of the mid-price process comes from relatively low frequencies (corresponding to a period longer than 8 hours). On the other hand, Figures 5 and 6 illustrate that, as the directional-change threshold increases, the spectral density tends to shift further towards the left-hand side of the graph, that is towards the lower frequencies. All the significant peaks correspond in fact to periods longer than

¹The graphs show only frequencies below the median value, as to focus on the significant peaks of the spectral densities.

100 and 340 hours, respectively for thresholds of 0.75% and 1.5%. This is perfectly consistent with the expectation that, as the directional-change threshold increases, the time needed to trigger that threshold also increases (for an equivalent result in the time domain see scaling law (5) in [Glattfelder, Dupuis, and Olsen, 2009](#), which relates the average time interval for a directional change of threshold Δx_{dc} to occur, to the size of the threshold).

In particular, the scaling property of directional-change durations assures that in the frequency domain, periodicities associated with a particular threshold are replicated in the empirical spectral densities of directional-change events of those of lower thresholds. This property is illustrated for the FX pair EUR-JPY in [Table 3](#), which shows for a subset of cases the periods expressed in hours, associated with significant peaks of the estimated spectral density at the 95% confidence level. It can be seen that the periodicities associated with the 150bp threshold (see [Figure 6](#)) are replicated in the spectral density of directional-change events sampled with both the 75bp and 50bp thresholds (see [Figure 5](#) and [4](#)). Similarly, the additional periodicities associated with the 75bp threshold are again propagated in the 50bp threshold and so on. If we observe a 150bp price change about every 360 hours on average (see [Table 3](#)), for example, we should also observe a 75bp *and* a 50bp price change at each full cycle.² In other words, for the same frequency in the spectral density of a directional change series sampled with a higher threshold, one can also find a corresponding significant peak in the spectral density sampled with a lower threshold. This remarkable behaviour in the frequency domain is due to the scaling

²Differences in the values might occur as “rounding errors” as the periods are actually calculated as reciprocal values from the original frequencies that are expressed in Hz. As discernible in [Table 3](#), the lower the reported period (i.e., the higher the estimated frequency), the smaller are the gaps between the values across the different thresholds.

Threshold		
150 bp	75 bp	50 bp
1153.5758	1184.5056	1187.6309
961.3131	987.0880	989.6924
823.9827	846.0754	848.3078
720.9849	740.3160	742.2693
	658.0586	659.7949
	538.4116	539.8322
	493.5440	494.8462
	394.8352	395.8770
360.4924	370.1580	371.1347
	348.3840	349.3032
	329.0293	329.8975
		312.5344
		296.9077
		247.4231
	236.9011	237.5262
		212.0769
	204.2251	204.7639
	197.4176	197.9385
	191.0493	191.5534
		185.5673
		179.9441
	174.1920	174.6516
	169.2151	169.6616
	164.5147	164.9487
		148.4539
	144.4519	144.8330

Table 3: Scaling property of directional-change durations for the currency pair EUR-JPY (time period: 11/2008 - 01/2009). The table shows, for a subsample of three different thresholds $\Delta x_{dc} = \{50bp, 75bp, 150bp\}$, the periods in hours associated with significant peaks of the estimated spectral density at 95% confidence level.

properties in the time domain, and is evident for other thresholds for the EUR-JPY FX rate, as well as for the other currency pairs investigated in this study.

In general, similar conclusions can be drawn analysing a different currency pair. For example, the spectral density of the AUD-HKD mid-price exhibits periodicities similar to those of the corresponding SDF of the currency pair EUR-JPY, showing an analogous trend, as the period associated with the directional changes tends to get longer as the directional-change threshold increases. Results for the other FX

pairs with similar implications are shown in the appendix (see Figures 7 to 16).

4 Concluding Remarks

UHF data are observed in real-time and therefore are characterised by the irregularity of time intervals between two consecutive events. This paper combines the LSFT and an event-based approach, to analyse foreign exchange tick-by-tick data. The Lomb-Scargle Fourier transform implicitly takes into account the non-periodic property of UHF data without the need to first transform the data to a periodic array. Using empirical transaction data from FX markets and adopting an event-based time scale (known as intrinsic time), the spectral analysis shows that various parts of the whole price process display different periodic patterns, revealed by the energy of the process in the respective frequency domain. The period associated with these patterns tends to increase as the directional-change threshold increases, confirming similar results in other studies (see e.g. [Glattfelder, Dupuis, and Olsen, 2009](#)).

Further, this framework can be generalised to a multivariate scheme (see e.g. [Schulz and Stattegger, 1997](#)), which allows the analysis of dependencies between different variables (e.g. price, volume, etc.) that can not be easily captured by current econometric models in the literature. On the contrary, comovements of multivariate time series can be detected in the frequency domain by computing the cross, coherency and phase spectra. This avenue of research represents a significant challenge and is an area for future investigation.

References

- BAUWENS, L., AND N. HAUTSCH (2009): “Modelling Financial High Frequency Data Using Point Processes,” in *Handbook of Financial Time Series*, ed. by T. G. Andersen, R. A. Davis, J. Kreiss, and T. Mikosch, pp. 953–979. Springer.
- BISIG, T., A. DUPUIS, V. IMPAGLIAZZO, AND R. OLSEN (2009): “The scale of market quakes,” Working paper, arXiv:0909.1690v1.
- BLOOMFIELD, P. (2000): *Fourier Analysis of Time Series: An Introduction*. Wiley-Interscience, 2 edn.
- BRETTTHORST, G. L. (2001): “Generalizing the Lomb-Scargle periodogram—the non-sinusoidal case,” in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, vol. 568, pp. 246–251.
- CHORDIA, T., R. ROLL, AND A. SUBRAHMANYAM (2001): “Market Liquidity and Trading Activity,” *The Journal of Finance*, 56(2), 501–530.
- COX, D., AND V. ISHAM (1980): *Point Processes*. Chapman & Hall/CRC.
- DACOROGNA, M. M., R. GENÇAY, U. A. MÜLLER, R. B. OLSEN, AND O. V. PICTET (2001): *An Introduction to High-Frequency Finance*. Academic Press.
- DEVILLIERS, V. (1933): *The Point and Figure Method of Anticipating Stock Prices: Complete Theory & Practice*. Windsor Books.
- ENGLE, R. F. (2000): “The Econometrics of Ultra-High-Frequency Data,” *Econometrica*, 68, 1–22.

- ENGLE, R. F., AND J. R. RUSSELL (1998): “Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data,” *Econometrica*, 66, 1127–1162.
- GIAMPAOLI, I., W. L. NG, AND N. CONSTANTINOU (2009): “Analysis of ultra-high-frequency financial data using advanced Fourier transforms,” *Finance Research Letters*, 6(1), 47–53.
- GLATTFELDER, J., A. DUPUIS, AND R. OLSEN (2009): “An extensive set of scaling laws and the FX coastline,” Working paper, arXiv:0809.1040.
- GRANGER, C. W. J., AND R. ENGLE (1984): “Applications of spectral analysis in econometrics,” in *Handbook of Statistics*, ed. by D. Brillinger, and P. R. Krishnaiah, vol. 3, chap. 5, pp. 93–109. Elsevier Science Publishers.
- GUILLAUME, D. M., M. M. DACOROGNA, R. R. DAVÉ, U. A. MÜLLER, R. B. OLSEN, AND O. V. PICTET (1997): “From the bird’s eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets,” *Finance and Stochastics*, 1(2), 95–129.
- HALL, A., AND N. HAUTSCH (2006): “Order aggressiveness and order book dynamics,” *Empirical Economics*, 30, 973–1005.
- HAUTSCH, N. (2004): *Modelling Irregularly Spaced Financial Data*. Springer, Heidelberg.
- HEINEN, A., AND E. RENGIFO (2007): “Multivariate autoregressive modeling of time series count data using copulas,” *Journal of Empirical Finance*, 14, 564–583.

- HORNE, J. H., AND S. L. BALIUNAS (1986): “A prescription for period analysis of unevenly sampled time series,” *Astrophysical Journal*, 302, 757–763.
- IACOBUCCI, A. (2005): “Spectral Analysis for Economic Time Series,” in *New Tools of Economic Dynamics*, ed. by J. Leskow, M. Puchet, and L. Punzo, Lecture notes in economics and mathematical systems, chap. 12, pp. 203–220. Springer, Berlin.
- LOMB, N. R. (1976): “Least-squares frequency analysis of unequally spaced data,” *Astrophysics and Space Science*, 39, 447–462.
- LYONS, R. K. (2001): *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge, Mass.
- MACDONALD, R. (2007): *Exchange Rate Economics: Theories and Evidence*. Routledge, London.
- MANDELBROT, B., AND H. M. TAYLOR (1967): “On the distribution of stock prices differences,” *Operations Research*, 15(6), 1057–1062.
- PACURAR, M. (2008): “Autoregressive Conditional Duration Models In Finance: A Survey Of The Theoretical And Empirical Literature,” *Journal of Economic Surveys*, 22(4), 711–751.
- PRESS, W. H., B. P. FLANNERY, S. A. TEUKOLSKY, AND W. T. VETTERLING (1992): *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2 edn.
- PRESS, W. H., AND G. B. RYBICKI (1989): “Fast algorithm for spectral analysis of unevenly sampled data,” *Astrophysical Journal*, 338, 277–280.

- PRIESTLEY, M. B. (1981): *Spectral Analysis and Time Series. Volume 1: Univariate Series*. Academic Press.
- RYDBERG, T. H., AND N. SHEPHARD (2003): “Dynamics of Trade-by-Trade Price Movements: Decomposition and Models,” *Journal of Financial Econometrics*, 1(1), 2–25.
- SCARGLE, J. D. (1982): “Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data,” *Astrophysical Journal*, 263, 835–853.
- SCHULZ, M., AND K. STATTEGGER (1997): “Spectrum: spectral analysis of unevenly spaced paleoclimatic time series,” *Computers & Geosciences*, 23, 929–945.
- VAN DONGEN, H., E. OLOFSEN, J. VAN HARTEVELT, AND E. KRUYT (1999): “A Procedure of Multiple Period Searching in Unequally Spaced Time-Series with the Lomb-Scargle Method,” *Biological Rhythm Research*, 30, 149–177.
- WARE, A. F. (1998): “Fast Approximate Fourier Transforms for Irregularly Spaced Data,” *SIAM Review*, 40, 838–856.
- WARNER, R. M. (1998): *Spectral Analysis of Time-Series Data*. The Guilford Press.
- WYCKOFF, R. D. (1910): *Studies in tape reading*. The Ticker Publishing Company, New York.

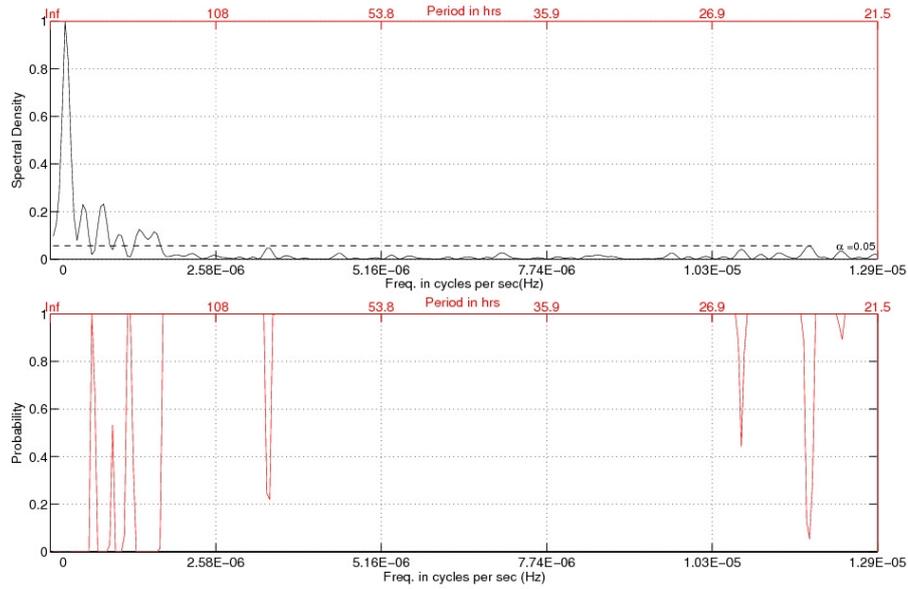


Figure 7: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 160$) of mid-price DC events of threshold 0.75% for the AUD-HKD FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

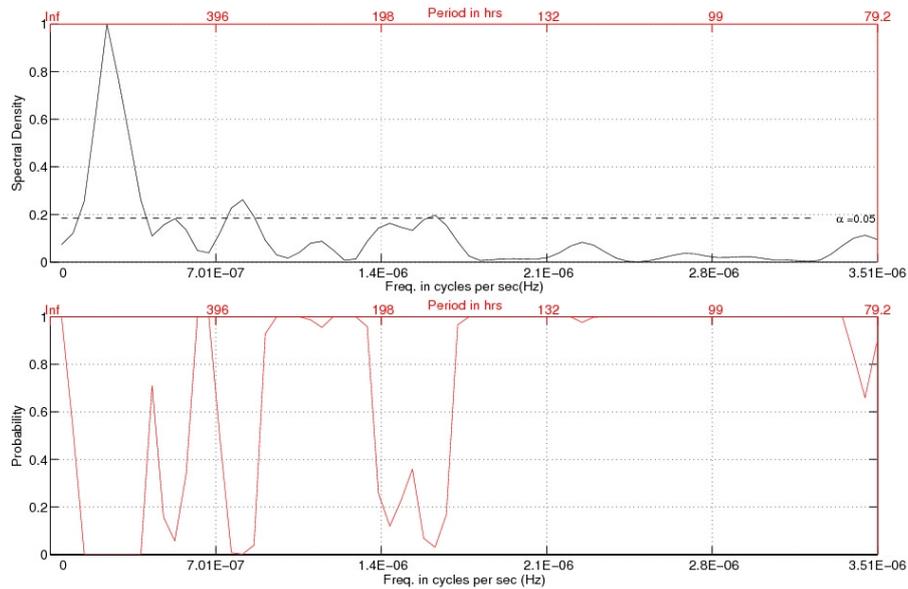


Figure 8: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 42$) of mid-price DC events of threshold 1.5% for the AUD-HKD FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

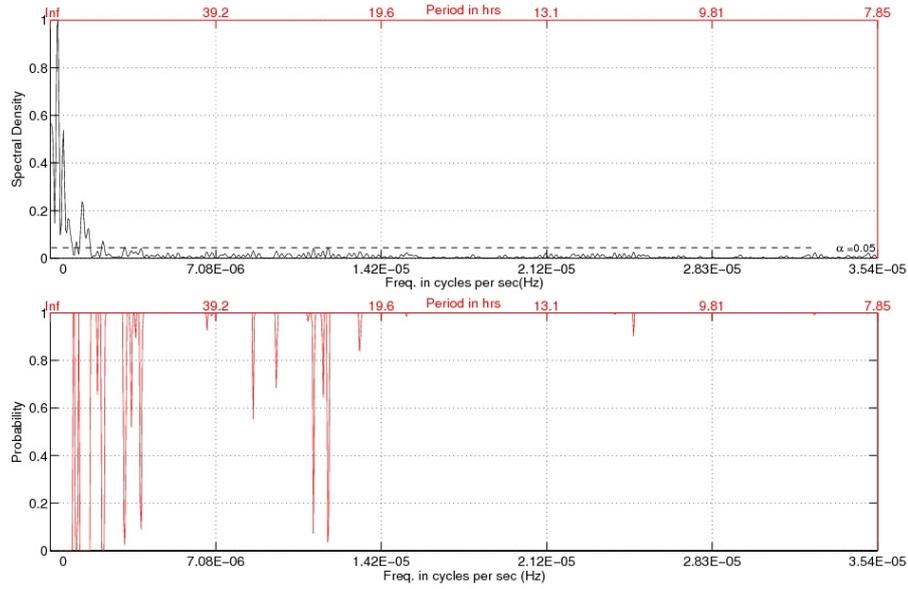


Figure 9: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 222$) of mid-price DC events of threshold 0.75% for the AUD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

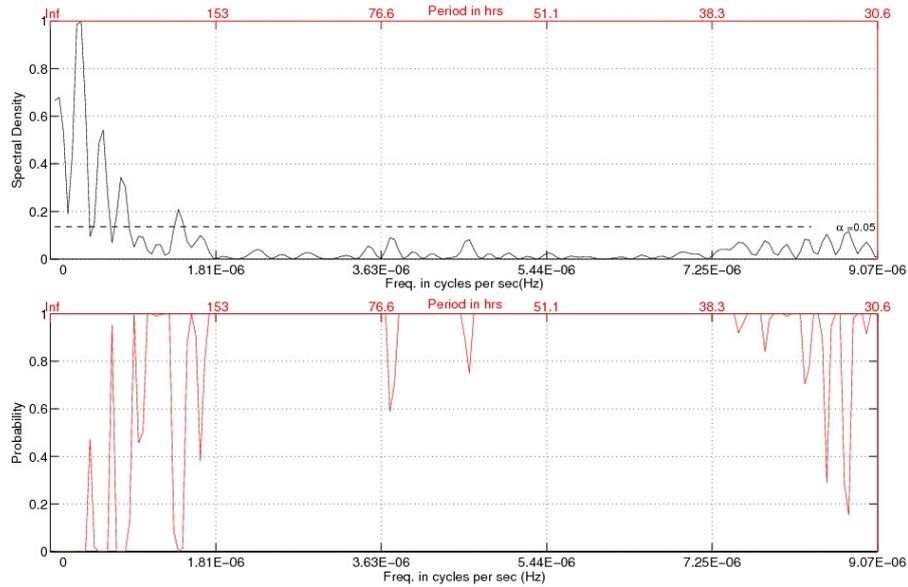


Figure 10: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 60$) of mid-price DC events of threshold 1.5% for the AUD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

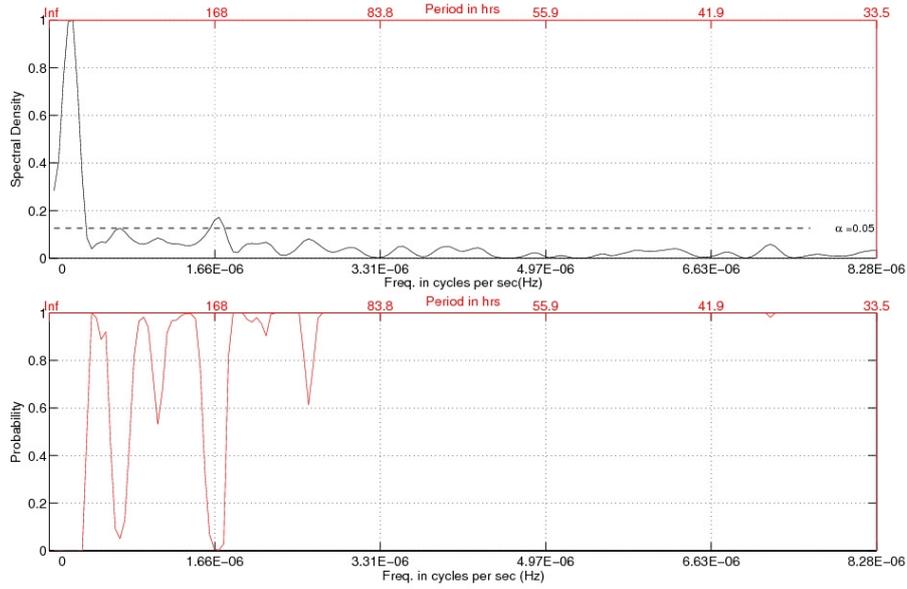


Figure 11: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 65$) of mid-price DC events of threshold 0.75% for the EUR-USD FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

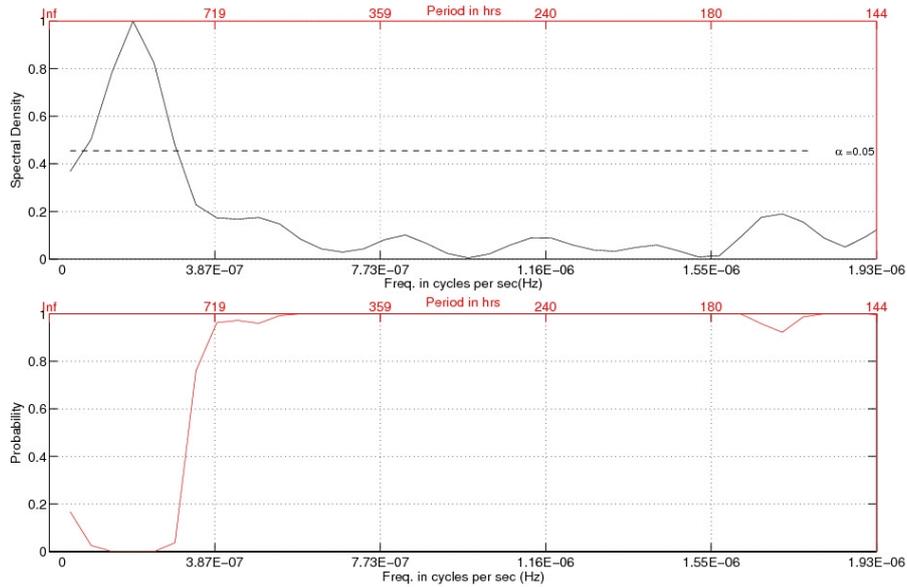


Figure 12: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 15$) of mid-price DC events of threshold 1.5% for the EUR-USD FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

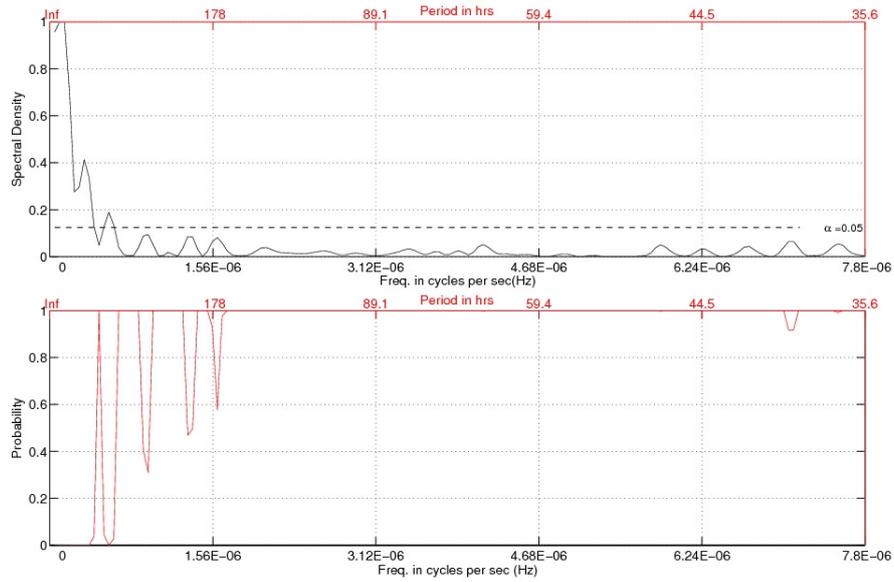


Figure 13: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 65$) of mid-price DC events of threshold 0.75% for the HKD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

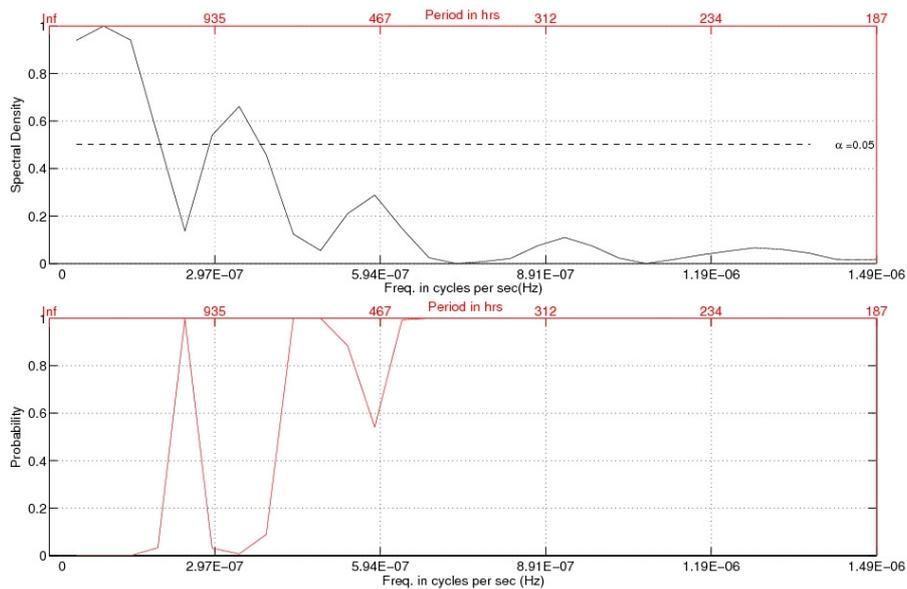


Figure 14: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 13$) of mid-price DC events of threshold 1.5% for the HKD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

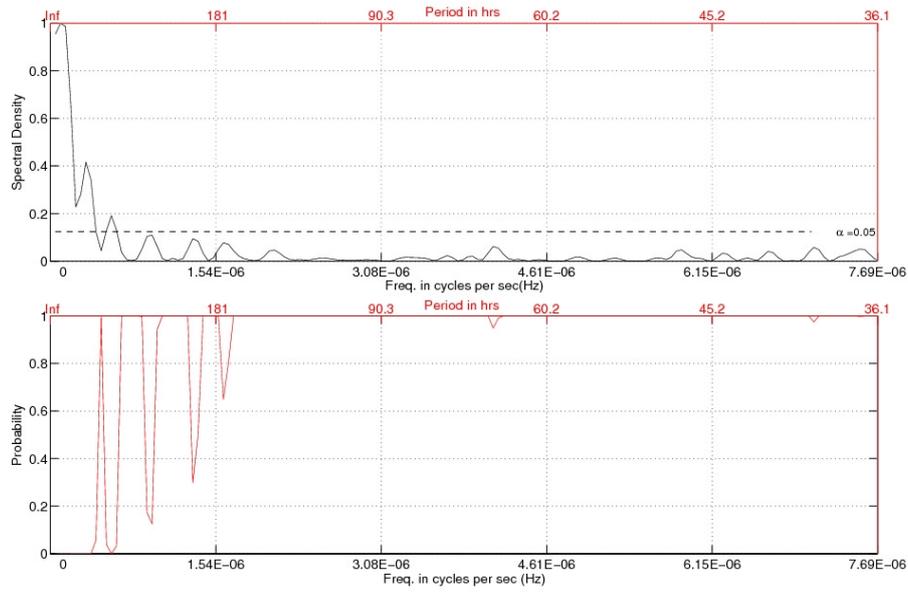


Figure 15: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 65$) of mid-price DC events of threshold 0.75% for the USD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.

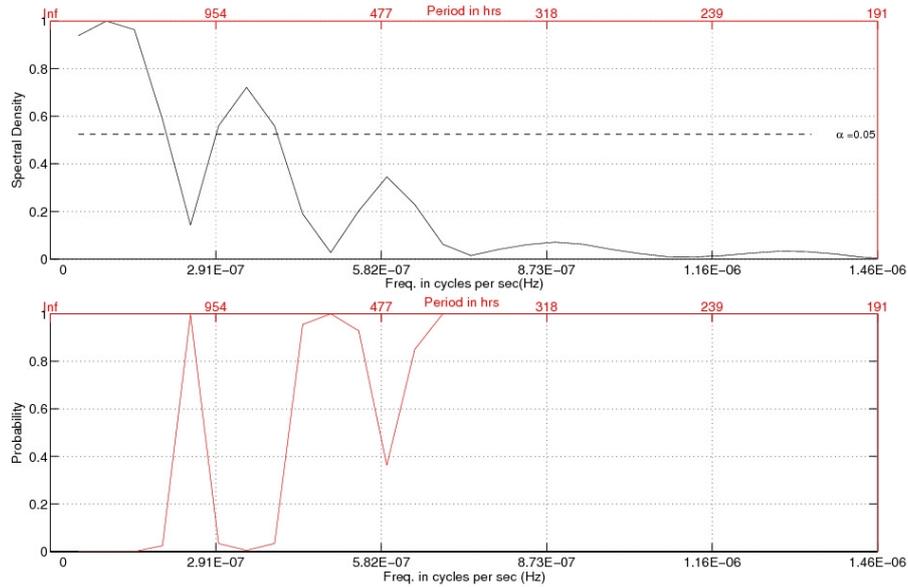


Figure 16: The upper panel shows the empirical spectral density (normalised by $\max(SDF_{LS}) \approx 12$) of mid-price DC events of threshold 1.5% for the USD-JPY FX rate. The lower panel shows the corresponding false alarm probabilities of the estimated spectral density at 95% s.l. for a given frequency expressed in Hz. The associated period in hours is indicated in the secondary x -axis.