

Stylized Facts of the FX Market

Transactions Data:

An Empirical Study

Monira Aloud¹, Maria Fasli², Edward Tsang,³
Alexandre Dupuis⁴ and Richard Olsen⁵

Abstract

In this paper, we focus on studying the statistical properties (stylized facts) of the transactions data in the Foreign Exchange (FX) market which is the most liquid financial market in the world. We use a unique high-frequency dataset of anonymised individual traders' historical transactions on an account level provided by OANDA. To the best of our knowledge, this dataset can be considered to be the biggest available high-frequency dataset of the FX market individual traders' historical transactions. The established stylized facts can be grouped under three main headings: scaling laws, seasonality statistics and correlation behaviour. Our work confirms established stylized facts in the literature

¹ College of Business Administration, King Saud University, KSA.

E-mail: mealoud@ksu.edu.sa

² School of Computer Science and Electronic Engineering, University of Essex, UK.

E-mail: mfasli@essex.ac.uk

³ School of Computer Science and Electronic Engineering, University of Essex, UK.

E-mail: edward@essex.ac.uk

⁴ Olsen Ltd., Seefeldstrasse 233, 8008 Zurich, Switzerland. E-mail: alex@olsen.ch

⁵ Olsen Ltd., Seefeldstrasse 233, 8008 Zurich, Switzerland. E-mail: richardo@olsen.ch

but also goes beyond those as we have discovered four new scaling laws and established six quantitative relationships amongst them, holding across EUR/USD and EUR/CHF transactions.

JEL classification numbers: G15

Keywords: Foreign Exchange (FX) market; stylized facts; high-frequency dataset; scaling laws; seasonality statistics.

1 Introduction

Financial markets are active and dynamic environments generating increasingly large volumes of data at frequencies higher than on a daily basis [1]. Such data have been the focus of study by both the academic community and industry analysts from a number of perspectives. Their study and analysis can reveal many properties of market behaviour, the strategic behaviour of market participants, the impact of participants' behaviour on the market, and the competition among related markets. Consequently, such studies improve our understanding of the different phenomena which emerge in financial markets [1].

One key area in the study of financial markets data is the establishment of their properties through statistical analysis [2]. These statistical properties are known as stylized facts [3]. In essence, stylized facts characterise markets and can provide useful insights into their workings, the behaviour of individual traders, the forces that drive behaviour and their impact on the market dynamics. Hence, their study is both imperative and fundamental in better understanding markets which is needed both for decision support and devising appropriate strategies. In addition, establishing stylized facts with regard to the behaviour of traders and their trading activities is an important step in modelling financial markets [1]. In building models of financial markets, stylized facts can be used as verification criteria [4] to confirm that an agent-based market (ABM) is indeed a model of the real market if it is able to reproduce the stylized facts to a satisfactory extent. A number of studies have established the stylized facts associated with price returns [1, 2, 3], order books [5, 6] and

transactions data [1, 7] in the market.

Despite the availability of studies establishing the stylized facts of price returns and order flow in the financial markets, the establishment of the stylized facts relating to transactions data in the high-frequency Foreign Exchange (FX) market is still in its infancy [1]. To establish the stylized facts of FX market traders' behaviour, we need to explore the high-frequency data (HFD) of their historical trading activities in the market. HFD allows us to explore some of the properties of trading behaviour that cannot be observed at lower frequencies (e.g. daily data) [1]. However, such HFD are very difficult to obtain.

In this paper, and motivated by the need to better characterise and understand the FX market, we present a set of stylized facts that have been identified through the study of a high-frequency transactions dataset. Our study involves a unique high-frequency dataset representing the physical transactions of anonymous accounts from January 1, 2007 to March 5, 2009. The large size of the dataset and, most importantly, the level of detail of the dataset, allows for a microscopic analysis of the traders' trading behaviour. In addition, our study makes use of EUR/USD and EUR/CHF bid and ask prices from January 1, 2007 to March 5, 2009. Our aim is to establish stylized facts of the transactions data in the high-frequency FX market in order to gain insights into the workings of the market, but also to establish a benchmark for validating FX agent-based market models in the future. We observe the behaviour of the FX market traders' and we establish stylised facts based on their collective behaviour which focus on: scaling laws, seasonality statistics and correlation behaviour. These independent stylized facts apply to transactions data and describe the trading activity in FX markets from different angles.

The organization of this paper is as follows. The following section provides a brief review of related work. Section 3 presents an overview of the FX market, provides a brief description of the datasets used in our study and the filtering process employed for the high-frequency transactions dataset. We present the study of the established stylized facts and the description and illustration of the new scaling laws and quantitative relations in section 4. The paper ends with the conclusions and avenues for future work.

2 Related Work

Identifying stylized facts in financial and other markets is an important research activity. In this section, we describe some of the work that has been done on establishing the stylized facts of financial assets and their returns, order flows, traders' behaviour and their collective transactions in the market.

Dacorogna et al. [1] provide a good overview of the main stylized facts with regard to foreign exchange rates. These stylized facts are grouped under four main headings for HFD: autocorrelation of return, distributional issues, scaling laws, and seasonality. In their work, they found remarkable similarity between the stylized facts of the different types of asset in financial markets. Cont [3] presents a review of the asset's price return stylized facts at low and high frequencies in various types of financial markets. These stylized facts of asset returns are: distributional properties, tail properties, and linear and nonlinear dependence of returns in time [3]. Several works have established the stylized facts of long memory volatility [8, 9, 10, 11, 12, 13] and fat tails for daily and intra-day data [14, 13].

A significant body of works has determined the scaling laws for a wide range of market data and time intervals [15, 16, 17, 2, 18, 11, 1, 19, 20]. There is one scaling law that is widely reported in the literature [15, 16, 17, 2, 18, 11, 1, 19, 20]: the size of the mean absolute change of the price is scaled to the size of the time interval of its occurrence. This scaling law has been applied to studying volatility and measuring risk [21, 22, 23, 24]. The scaling law discovered by Guillaume et al. relates the number of so-called directional changes to the directional-changes sizes [2]. Glattfelder et al. discovered 12 independent new scaling laws in foreign exchange data series holding across 13 currency exchange rates [20]. Their statistical analysis depends on the so-called directional-change event approach. The discovered scaling laws give an estimation of the length of the price-curve coastline which turns out to be long. The first scaling law discovered in [20] relates the average number of ticks observed during a price move of size Δx to the size of that threshold Δx . According to Glattfelder et al., a tick is defined as a price move larger than 0.02%. The second scaling law counts the average yearly number of price moves of size Δx . The third scaling law relates the average difference between the high and low price levels during a time interval Δt to the size of that time interval Δt . Law four relates

the average time interval for a price change of size Δx to occur to the size of the threshold and similarly law five considers directional changes instead of a price move of size Δx . A set of six scaling laws emerge from the so-called total-move of the price which decompose into directional change and overshoot events. The last scaling law considers cumulative price moves for a price move of size Δx to this threshold Δx . These scaling laws provide a foundation for better understanding the foreign exchange market. The directional change event approach used in [20] is explained in section 4.1.

A number of works have analysed and defined the effect of order flow using analytical models [5, 6, 25, 26, 27]. Although there are researchers who have studied order flow in the FX market, they have acquired either large samples of low frequency datasets [6], or small samples of high-frequency datasets [25], neither of which are at an account level.

Researchers have also used psychology-based approaches to study how the traders' psychological makeup impacts on their trading decisions in terms of price changes and new events in the market. These works relate to traders' herding behaviour [28], feedback trading in the market [29, 30, 31] and traders' heterogeneous expectations and beliefs [32, 33, 34, 35, 36, 37].

However, the establishment of stylized facts of the FX market transactions data, is still in its infancy due to the limited availability of high-frequency data of market transactions. Dacorogna et al. [1] established stylized facts with regard to the seasonality of transactions in the FX markets by quantifying the trade frequency and volume using price tick data. They show that the intraday dynamics of transactions in the FX market exhibits a double U-shape or camel-shape pattern. Ito et al. [7] established stylized facts associated with seasonality and correlation behaviour for the USD/YEN and the EUR/USD. Their work confirmed the existence of the double U-shape pattern of intraday transactions for Tokyo and London participants. They have also found that the price changes and the trade volumes have a positive correlation.

In this work, we aim to confirm and extend the seasonality and correlation statistics work described in [1, 7]. Also, we aim to confirm and extend the scaling laws discovered in [20]. These stylized facts have the potential to improve our understanding of the dynamic behaviour of FX markets and can help us explain the emergence of different patterns and phenomena. In addition, they can be valuable tools for forecasting and decision support systems and

modelling strategies.

3 The FX Market and Datasets

The FX market is where the buying, selling and exchanging of currencies takes place and it is considered the largest, most liquid and most efficient financial market in the world [1]. As such it is not a single market, but it is composed of a global network of FX markets that connect investors from all around the world. It is also a decentralized market as there is no central marketplace and transactions are conducted over the counter. Furthermore, the FX market has no business hour limitations and operates 24 hours a day, 7 days a week. Traders can be governments, central banks, commercial banks, retail investors, institutional investors, etc.

With the advent of retail market-maker FX market online platforms, individual retail traders represent an important element of growth for the FX market. Most FX trading firms are market-makers. A market-maker is a company which provides liquidity for a particular currency pair, and quotes both a buy and a sell price for such a currency pair on its platform. The market-maker buys from and sells to its clients and other market-makers, to make a profit on the bid-offer spread or return. In other words, the market-maker takes the opposite side of a trade and earns its commission from the difference between the bid and the offer price.

3.1 The Datasets

As the FX market is the largest and most liquid financial market in the world, this makes it an important source of high-frequency data (HFD) [1]. HFD represents an extremely large amount of data recorded at frequencies higher than on a daily basis. HFD have unique features that are lacking in data recorded at lower frequencies, such as intra-day data. These HFD are irregularly spaced in time whereas low frequency data are regularly spaced in time [1, 2]. Using HFD is fundamental to the understanding of financial markets since participants determine their trading decisions by observing HFD [1].

In this study, we used two high-frequency historical datasets from OANDA Corporation, short for Olsen And Associates. OANDA is an online market-maker trading platform for the trading of foreign currencies and it serves a variety of traders, from individual retail traders to corporations and financial institutions. In OANDA, traders trade under the same terms and conditions, particularly under the same prices with competitive spreads. It is worth highlighting, that there is no trading platform except for OANDA that stores the details of the historical transactions over a long time horizon.

The first dataset represents 2.25 years data samples of high-frequency EUR/USD and EUR/CHF bid and ask prices. The sample range is from January, 1 2007 to March, 5 2009. Each data record contains three fields: (a) a bid, and (b) an ask price at (c) a timestamp. Throughout the paper, the following definition of mid-price is used:

$$p_{m,t} = \frac{(b_t + a_t)}{2} \quad (1)$$

where $p_{m,t}$ is the mid-price of a currency pair at time t , and b_t and a_t are the bid and ask price respectively at time t .

The second dataset represents a unique high-frequency dataset of individual traders' historical transactions at an account level made available on an anonymous basis spanning 2.25 years, from January, 1 2007 to March, 5 2009. The dataset includes about 147 million transactions carried out by 45,845 different accounts trading in 48 different currency pairs under the same terms and conditions. Each transaction includes: the transaction type, the transaction timestamp, the traded currency pair, the execution/transaction price, the units and the amount traded. For further information on the datasets, we refer the interested reader to [38].

Although the transaction dataset includes transactions in 48 different currency pairs, the scaling laws analysis and results only apply to two of them (EUR/USD and EUR/CHF). This is because the analysis of scaling laws depends on the availability of high-frequency price datasets. In this study, we acquired only two datasets of prices - EUR/USD prices and EUR/CHF prices. The correlation analysis and results apply only to EUR/USD transactions and prices due to the similarity of reporting the analysis of the EUR/CHF results. On the other hand, the seasonality analysis and results include transactions in all of the 48 different currency pairs.

3.2 Filtering the Datasets

The use of HFD comes with a set of challenges as such datasets may possibly contain observations that are not consistent and compatible in terms of actual market activity. Hence before analysing such datasets there is a need to process and clean them. Erroneous and misleading observations may possibly be the result of the institution's internal system storage procedures, which may entail using dummy ticks [1]. Data gaps in HFD may possibly result from computer system errors during the process of storing the data [39, 40]. Therefore, a clean dataset is an essential pre-condition for the analysis phase of the HFD. Failure to recognize erroneous and misleading data may possibly cause ambiguous results in the statistical analysis. The procedure for filtering HFD depends on the structure of the HFD and the types of error. In the literature, a variety of customized approaches have been adopted for filtering HFD [41, 40, 1, 42].

The HF transactions dataset used in this study was filtered to remove any erroneous and misleading transactions in terms of the individual trader's actual trading activity. The major issues of the HF transactions dataset reside in: (a) OANDA's system storage procedure, involving storing one trade (an executed order) in several transactions, (b) OANDA's internal interest payment procedure producing dummy transactions, and (c) an unexpected sudden drop of the flow of number of transactions with an increasing number of accounts. We validate the reliability and consistency of the filtered dataset by tracking for each of the account's traded currency pairs, each transaction's traded units bought and sold in sequence in terms of the transactions' execution time.

By carrying out the filtering procedure, we have confirmed a clean dataset, reduced to a total of ~ 59 million transactions from ~ 147 million transactions. An important step in the filtering process is the validation of the clean dataset. For a detailed description of the filtering procedure of the OANDA HF transactions dataset and its validation, we refer interested readers to [38].

4 Stylized Facts

In this section, we describe the results of our study regarding the stylized facts that can be observed in the HF dataset. The results include the establishment of four new empirical scaling laws.

4.1 Scaling Laws

Scaling laws describe the average absolute price returns and the average market transactions as functions of their time intervals over which they are measured. The time intervals vary from a few seconds to one or more days. These scaling laws are proportional to a power of the time interval size. A scaling law relation shows a simple functional relationship between the occurrences of an observed statistical property measured at different time intervals. It gives a direct relationship between average price movements and the average number of transactions, and their volumes, measured at different time intervals. The discovery of scaling laws in FX market data reveals and explains the patterns and mechanisms that exist in the FX market. Following Glattfelder et al.'s work in [20], we have extended the set of stylized facts of the FX market by observing four new scaling laws, next to establishing six quantitative relationships amongst them, holding across EUR/USD and EUR/CHF transactions. Since prices in the market change at uneven time intervals, the measurement of market trading activity needs to be adaptive beyond the notion of physical time scale changes. In this regard, our statistical analysis depends on an event-driven approach – the so-called directional-change event approach. The directional-change event approach characterizes price movements in the price time series where any occurrence of a directional-change (DC) event represents a new intrinsic time unit, independent of the notion of physical time changes. Prior to the study of the price time series, two variables are defined: the last high and low prices which are set to the initial price at the start of the price's sequence. Given a threshold of size Δx , a DC event is a price change of size Δx from the last high or low price whether it is a downturn or an upturn event, respectively. A DC event of size Δx is usually not followed by an opposite DC event but by an overshoot (OS) event [20]. An OS event is the excess price move from one DC event of size Δx to the next DC event. In other words, an OS event is defined as the difference between the price at which the last DC event occurred, and the next extrema. The extrema is the last high price when the previous DC event is an upturn event, or the last low price when the previous DC event is a downturn event [43]. The last high and last low prices are afterwards reset to the current market price at the time a DC event occurs [43]. For the period of an upward trend, the last high price of an asset is continuously adjusted to the maximum of the asset's (a) current price and (b)

last high price. Conversely, for the period of a downward trend, the last low price of asset is continuously updated to the minimum of the asset's (a) current price and (b) last low price. For more details regarding the directional-change event approach we refer the interested reader to [44].

4.1.1 Empirical Evidence on Existing Stylized Facts

In this section, we have empirically confirmed from the existing literature four scaling laws and three quantitative relationships amongst them. These four scaling laws hold across EUR/USD and EUR/CHF prices. Law (a) is observed in a physical time scale whereas the other laws are observed in an intrinsic time scale (i.e. using the directional-change event approach). The computation of laws (b), (c) and (d) relies on the detection of DC and OS events instead of focusing on the stochastic nature of the data-series.

The scaling law (a) was discovered by Miller et al. in [15] and relates the size of the average absolute mid-price change (return), sampled at time intervals Δt , to the size of the time interval

$$\langle |\Delta x_t| \rangle = \left(\frac{\Delta t}{C_x} \right)^{E_x} \quad (2)$$

where C_x and E_x are the scaling law parameters and a price move Δx_t at time t is defined as

$$\Delta x_t = \frac{(p_{m,t} - p_{m,t-1})}{p_{m,t-1}} \quad (3)$$

The scaling laws parameters are the results of the line fit, where E_x is the slope and C_x is the intercept. The slope measures the proportional change of the average absolute mid-price change due to an increment in the time interval.

Using the directional-change event approach as explained in [44], we define DC and OS events in EUR/USD and EUR/CHF mid-price time series for a set of thresholds of different size, ranging from 0.10% to 0.80%. We have confirmed the scaling law (b) which was discovered by Guillaume et al. [2]. Law (b) relates the number $N(\Delta x_{DC})$ of DC events to the size of the DC event Δx_{DC}

$$N(\Delta x_{DC}) = \left(\frac{\Delta x_{DC}}{C_{N,DC}} \right)^{E_{N,DC}} \quad (4)$$

where $C_{N,DC}$ and $E_{N,DC}$ are the scaling law parameters.

The analysis of the price data suggests that the two scaling laws are exhibited in the data as discovered by Glattfelder et al. in [20]:

1. Law (c) relates the time during which events occur to the size of these events. Given a fixed percentage threshold, the average time interval $\langle \Delta t_x \rangle$ for a price move of size Δx to occur is scale-invariant to the size of the threshold

$$\langle \Delta t_x \rangle = \left(\frac{\Delta x}{C_{t,x}} \right)^{E_{t,x}} \quad (5)$$

where $C_{t,x}$ and $E_{t,x}$ are the scaling law parameters.

2. Law (d) counts the average number of ticks observed during every event. Given a fixed percentage threshold, the average number of ticks $\langle N(\Delta x_{tick}) \rangle$ observed during a price move of size Δx is scale-invariant to the size of this threshold

$$\langle N(\Delta x_{tick}) \rangle = \left(\frac{\Delta x}{C_{N,tick}} \right)^{E_{N,tick}} \quad (6)$$

where a tick is defined as an individual quote of bid and ask price by a market-maker, $C_{N,tick}$ and $E_{N,tick}$ are the scaling law parameters.

The scaling law parameters are estimated using a simple linear regression to model the relationship between a scalar dependent variable and one explanatory variable. We used the least square method to fit a regression line to the observed data by minimising the total of the squares of the vertical deviation in terms of each data point to the fitted line. When the observed data points lie on the fitted line, then the vertical deviation of the data points is zero. In the least square method, the adjusted R^2 value measures how accurately the line fit is in explaining the variation of the observed data. The adjusted R^2 value can be any value from 0 to 1, with a value closer to 1 signifying a better fit. The standard error of the fitted regression line measures the accuracy with which the regression line is measured. Laws (b), (c) and (d) are plotted in Figure 1. Table 1 reports the adjusted R^2 values and the standard error of the fit.

The average obtained results (given in Figure 1) of the different event thresholds that we considered confirm the three quantitative relationships among laws (c) and (d) which were discovered by Glattfelder et al. in [20]:

	EUR/USD				EUR/CHF			
	DC		OS		DC		OS	
	Adj. R ²	SE	AAdj. R ²	SE	Adj. R ²	SE	Adj. R ²	SE
Directional changes	0.99	0.35	-	-	0.99	0.44	-	-
Time interval	0.99	0.34	0.99	0.22	0.99	0.04	0.99	0.05
Tick numbers	0.99	0.36	0.99	0.35	0.99	0.43	0.99	0.45

Table 1: Estimated regression parameters: the adjusted R² values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events. The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009.

1. A DC event of size Δx_{DC} is followed by one OS event of the same size

$$\langle \langle \Delta x_{DC} \rangle \rangle \approx \langle \langle \Delta x_{OS} \rangle \rangle \quad (7)$$

2. An OS event takes twice as long as a DC event to unfold

$$\langle \langle |\Delta t_{OS}| \rangle \rangle \approx 2 \langle \langle |\Delta t_{DC}| \rangle \rangle \quad (8)$$

where $\langle \langle |\Delta t_{OS}| \rangle \rangle$ is the average time it takes an OS event to unfold while $\langle \langle |\Delta t_{DC}| \rangle \rangle$ is the average time it takes a DC event to unfold.

3. An OS event contains twice as many ticks as a DC event

$$\langle \langle N(\Delta x_{OS,tick}) \rangle \rangle \approx 2 \langle \langle N(\Delta x_{DC,tick}) \rangle \rangle \quad (9)$$

where $\langle \langle N(\Delta x_{OS,tick}) \rangle \rangle$ and $\langle \langle N(\Delta x_{DC,tick}) \rangle \rangle$ are the average tick numbers in an OS and a DC event, respectively.

4.1.2 The New Scaling Laws

As already demonstrated in the analysis in the previous section (4.1.1), the empirical evidence indicates a scaling behaviour observed in price data. Extending Glattfelder et al.'s work [20], we identified four new scaling laws and cross-checked our results by establishing six quantitative relationships amongst them, holding across EUR/USD and EUR/CHF transactions. It is important to highlight that the scaling laws reported in section 4.1.1 apply to price data,

while the scaling laws reported in this section apply to transactions data. A transaction represents an executable order in the market.

Scaling Law (1): Transaction Numbers

Given a fixed percentage threshold, the average number of transactions $\langle N(\Delta x_{trade}) \rangle$ observed during an event is scale-invariant to the size of this threshold

$$\langle N(\Delta x_{trade}) \rangle = \left(\frac{\Delta x}{C_{N,trade}} \right)^{E_{N,trade}} \quad (10)$$

where a transaction is defined as an executable order in the market, $C_{N,trade}$ and $E_{N,trade}$ are the scaling law parameters. In essence, this law counts the average number of transactions observed during every event. Law (1) is plotted in Figure 2 and Table 2 reports the adjusted R^2 values and the standard error of the fit. The results of Law (1) show that on average, an OS event contains roughly twice as many transactions as a DC event

$$\langle N(\Delta x_{OS,trade}) \rangle \approx 2 \langle N(\Delta x_{DC,trade}) \rangle \quad (11)$$

where $\langle N(\Delta x_{OS,trade}) \rangle$ and $\langle N(\Delta x_{DC,trade}) \rangle$ are the average transaction numbers in an OS and a DC event, respectively.

Scaling Law (2): Transaction Volumes

Given a fixed percentage threshold, the average volume of transactions $\langle V(\Delta x_{trade}) \rangle$ observed during an event is scale-invariant to the size of this threshold

$$\langle V(\Delta x_{trade}) \rangle = \left(\frac{\Delta x}{C_{V,trade}} \right)^{E_{V,trade}} \quad (12)$$

where a transaction volume for a single transaction i , denoted by V_i , is defined as the execution price times the number of currency units of the transaction. $C_{V,trade}$ and $E_{V,trade}$ are the scaling law parameters. In detail, this law counts the average volume of transactions observed during every event. Law (2) is plotted in Figure 3. Table 2 reports the adjusted R^2 values and the standard error of the fit. The results of the different event thresholds that we considered

show that on average, an OS event contains roughly twice as many volumes as a DC event

$$\langle V(\Delta x_{OS,trade}) \rangle \approx 2 \langle V(\Delta x_{DC,trade}) \rangle \quad (13)$$

where $\langle V(\Delta x_{OS,trade}) \rangle$ and $\langle V(\Delta x_{DC,trade}) \rangle$ are the average transaction volumes in an OS and a DC event, respectively.

Scaling Law (3): Number of Opening Positions

A position is opened once a trader has bought or short-sold any quantity of currency units. A position can be of type long (bought) or short (sold). An open position is closed by placing a transaction that has an equal amount and takes the opposite type to the open position.

Law 3 counts the average number of opening positions observed during every event of size Δx . Given a fixed percentage threshold, the average number of opening positions $\langle N(\Delta x_{OP}) \rangle$ observed during an event is scale-invariant to the size of this threshold

$$\langle N(\Delta x_{OP}) \rangle = \left(\frac{\Delta x}{C_{N,OP}} \right)^{E_{N,OP}} \quad (14)$$

where $C_{N,OP}$ and $E_{N,OP}$ are the scaling law parameters. Law (3) is plotted in Figure 4 and Table 2 reports the adjusted R^2 values and the standard error of the fit.

Scaling Law (4): Number of Closing Positions

This law counts the average number of closing positions observed during every event. Given a fixed percentage threshold, the average number of closing positions $\langle N(\Delta x_{CP}) \rangle$ observed during an event is scale-invariant to the size of this threshold

$$\langle N(\Delta x_{CP}) \rangle = \left(\frac{\Delta x}{C_{N,CP}} \right)^{E_{N,CP}} \quad (15)$$

where $C_{N,CP}$ and $E_{N,CP}$ are the scaling law parameters. Law (4) is plotted in Figure 5 and Table 2 reports the adjusted R^2 values and the standard error of the fit. The average results of Laws (3) and (4) show two important features:

	EUR/USD				EUR/CHF			
	DC		OS		DC		OS	
	Adj. R ²	SE						
Trade numbers	0.90	0.37	0.90	0.35	0.95	0.35	0.95	0.35
Trade volumes	0.91	0.38	0.90	0.35	0.99	0.35	0.95	0.36
Opening positions	0.90	0.35	0.91	0.35	0.94	0.34	0.94	0.36
Closing positions	0.90	0.36	0.90	0.35	0.94	0.34	0.94	0.36

Table 2: Estimated regression parameters for the new scaling laws: the adjusted R² values of the fits, plus their standard errors (SE), under DC and OS events. The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009.

1. An OS event contains roughly twice as many numbers of opening and closing positions as a DC event

$$\langle N(\Delta x_{OS,OP}) \rangle \approx 2 \langle N(\Delta x_{DC,OP}) \rangle \quad (16)$$

$$\langle N(\Delta x_{OS,CP}) \rangle \approx 2 \langle N(\Delta x_{DC,CP}) \rangle \quad (17)$$

where $\langle N(\Delta x_{OS,OP}) \rangle$ and $\langle N(\Delta x_{DC,OP}) \rangle$ are the average number of opening positions in an OS and a DC event, respectively. $\langle N(\Delta x_{OS,CP}) \rangle$ and $\langle N(\Delta x_{DC,CP}) \rangle$ are the average number of closing positions in an OS and a DC event, respectively.

2. An OS event contains roughly the same numbers of opening and closing positions. A similar feature holds for a DC event.

$$\langle N(\Delta x_{OS,OP}) \rangle \approx \langle N(\Delta x_{OS,CP}) \rangle \quad (18)$$

$$\langle N(\Delta x_{DC,OP}) \rangle \approx \langle N(\Delta x_{DC,CP}) \rangle \quad (19)$$

4.1.3 Discussion

Using an event-driven approach similarly to [20], we have extended the set of FX market stylized facts by discovering four new scaling laws relating to

the FX market transactions data. We cross-checked them and established six quantitative relations amongst them. The new scaling laws have the potential to enhance our understanding of the dynamic behaviour of FX markets as they describe the character and relationships of scaling patterns with regard to financial market transactions data.

From the results obtained, the transaction activity for EUR/USD and EUR/CHF currency pairs exhibit similar average behaviour. In relation to the quantitative relationships amongst the new scaling laws, it is clear, from the analysis, that an OS event contains approximately twice as many transactions, trading volumes and frequencies of opening/closing positions as a DC event. The explanation of why an OS event contains more transactions than a DC event relies on the fact that a DC event depicts a pattern set up for the move. Such a pattern influences traders' decisions, assuming that the current trend will continue in the same direction. An OS event, whether upward or downward, appears to indicate tendency amongst traders. A possible interpretation is that the traders become influenced by observing the direction of the overshoot price, trending upwards or downwards, and, essentially, go with the flow. On the contrary, what appears to happen during a DC event is that traders seem to hesitate and are more cautious in the frequency and volume of their transactions as the price trends upwards or downwards. In particular, it appears that a trader's response is slower to a small change in the direction of price movement than to a bigger change. Such trading behaviour is interpreted, in financial markets, as trend-following behaviour. In the literature, relating to financial markets, there are a considerable number of works that have established and studied the efficiency of trend-following strategies, such as [45, 46, 47, 48].

Understanding the dynamics of financial market behaviour is important as it can dramatically change the way in which traders observe and examine the markets. Using a framework of scaling laws properties enables us to develop forecasting models as regards the behaviour of prices, as well as the behaviour of market orders and transactions. For instance, a forecasting model uses a framework of scaling laws as a reference for predicting likely forthcoming price peaks or troughs as a means of aiding investment opportunities. Scaling laws can be used as tools for developing investment strategies in which orders are triggered at different time scales of price evolution. The scaling laws can be

used as a framework of signs to allow us to measure and characterise the level and flow of transactions and orders in the market. Such measurements and characteristics feed back into the investment strategies to enable such strategies to adapt to changing market behaviour. A systematic risk management methodology can be developed based on the dynamic framework of the statistical properties derived from the scaling laws of financial market data.

4.2 Seasonality

A periodic pattern exhibited in the market data is referred to as seasonal [1]. Seasonality in terms of price changes has been found in FX market data using intraday and intraweek frequencies [1, 7]. The analysis of intraday and intraweek seasonality has the advantage of operating from a very simple and clear definition. Such analysis relates quantities of different market activities to the time of the day or the week when these quantities are observed [1]. Thus, seasonality shows the average quantities observed for every hour of the day or of the week. The seasonality statistics reveal a great amount of information about the FX market's active trading hours, volume of transactions, and how the transactions flow might develop.

The seasonal patterns of transactions are associated with the inherent structure of the worldwide main market centres (e.g., London, Tokyo and New York) [1], and more specifically, their opening and closing times. Although technically the FX market operates continuously, it can be divided into three major trading sessions during which the volume of transactions peaks in a day: the East Asia, Europe and America [1]. Usually these three trading sessions are referred to as the Tokyo, London and New York trading sessions. Table 3 presents opening and closing hours in GMT of the major FX market trading sessions. When the trading hours of these market centres overlap, inevitably, the volume of the FX market transactions increases due to more traders participating in the market [1].

In this section, we study the intraday and intraweek seasonality of roughly 59 million transactions carried out by more than 40,000 individual accounts on the OANDA FX trading platform over 2.25 years.

Session	Major Market	Opening time	closing time
Asia	Tokyo	00:00:00	08:00:00
Europe	London	08:00:00	16:00:00
America	New York	13:00:00	21:00:00
Australia	Sydney	22:00:00	06:00:00

Table 3: FX market opening and closing business time for different geographical markets. The time zone is GMT. The markets are listed in the order of their opening times in GMT.

4.2.1 Intraday Seasonality

The intraday seasonality of transactions data is defined by means of the daily hourly changes in market transactions over a defined period of time [1]. For instance, the intraday seasonality of transaction numbers stands for the aggregated transactions occurring in each hour (hour 0 to hour 23) of the day, divided by the total transaction numbers in the underlying period. Figures 6 and 7 construct uniform time grids with 24 hourly intervals. Figure 6 shows the intraday seasonality of the transaction numbers made in terms of 48 different currency pairs. In contrast, Figure 7 shows the intraday seasonality for the EUR/USD (a) transaction numbers, (b) transaction volumes, (c) number of opening positions and (d) number of closing positions. We can spot from Figures 6 and 7 that the FX market activity exhibits a double U-shape or camel-shape pattern due to the different trading hours of FX market traders. This result is in line with the reported results of studies in the literature [1, 7]. The analysis of the intraday seasonality of FX market transactions data indicates the following:

- The transactions start to peak within the opening trading times of the market centres in the morning.
- The transactions decline during the lunch break of the market centres and then peak in the afternoon again.
- During the market centres' closing hours, the transactions gradually decline.
- The intraday seasonality of transactions has two peaks with the second peak being higher than the first. The first peak takes place when the

London trading session is open, while the second occurs when the London and New York trading sessions are open simultaneously.

- The lowest hourly transactions of the FX market outside weekends occurs during the first opening hour of the Sydney market, about 10:00 pm GMT time, when it is night time for the London and New York trading sessions.

The intraday statistics results do not differ (a) for traders through different currency pairs and with regard to the same currency pair (shown in Figure 6 and 7), and (b) over long and short time horizons.

4.2.2 Intra-week Seasonality

Due to the small number of participants in the FX market during weekends, weekends witness extremely low transactions [1]. This in turn implies weekly periodicity patterns. Consequently, it is important to add intra-week statistics to the intraday statistics when analysing the flow of FX market transactions [1].

The intra-week analysis of FX market transactions uses a uniform time grid of 168 hours from Monday 0:00 - 1:00 to Sunday 23:00 - 24:00 (GMT) to display the aggregated market transactions in each hour of the week [1]. This demonstrates the active trading period of the main FX market centres during the day, the same as for the intraday seasonality of market transactions.

Figure 8 shows the intra-week seasonality of the transaction numbers for 48 different currency pairs. Figure 9 shows the intra-week seasonality for the EUR/USD (a) transaction numbers, (b) transaction volumes, (c) number of opening positions and (d) number of closing positions. The patterns in Figures 8 and 9 can be explained by considering the weekend effect and the intraday seasonality (section 4.2.1). An extremely low level of transactions takes place from Friday at 21:00 to Sunday evening at 22:00 as a result of the weekend effect. In contrast, high level of transactions takes place as would be expected during the working days.

Similarly to the intraday statistics, the intra-week statistics do not differ (a) for transactions through different currency pairs and with regard to the same

currency pair (shown in Figures 8 and 9), and (b) over long and short time horizons.

4.3 Correlation Behaviour

A convenient way to discover statistical properties of market transactions data is to conduct a correlation study [1]. The linear correlation coefficient is a measurement of the level of the dependence between variables. It measures the relationship between two series of variables x and y , and is defined as follows:

$$\text{corr}(x_i, y_i) = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2 \sum_{i=1}^n (y_i - \bar{y})^2}} \quad (20)$$

with the sample means:

$$\bar{x} = \sum_{i=1}^n \frac{x_i}{n} \quad (21)$$

$$\bar{y} = \sum_{i=1}^n \frac{y_i}{n} \quad (22)$$

The correlation between quantities of financial market data reveals, at the same time, serial dependency. For the stock market, it has been confirmed by a number of studies in the literature that intraday price volatilities are correlated with the market activities in terms of the transaction data [49]. In contrast, for the FX market there are relatively few studies on the intraday behaviour of market transactions due to the difficulty of obtaining high-frequency data. However, Dacorogna et al. in [1] found across different currency pairs that the intraday volatilities are correlated with the activities measured in terms of the tick numbers. Ito et al. in [7] found that the price changes and the trade volumes have a positive correlation for the USD/YEN and the EUR/USD currency pairs they considered. In this section, we present an analysis of the correlation function of EUR/USD trading activity and mid-price volatility. This is done using hourly correlation calculation intervals over January, 1 2007 to January, 31 2009.

Figure 10 shows the correlation of EUR/USD hourly (a) trade numbers and volumes, (b) numbers of buy and sell executed orders, and (c) numbers of opening and closing positions. We can see that the plots in Figures 10 form an almost straight line sloping upwards to the right. These indicate a

	Correlation coefficient
Transaction numbers and volumes	+ 0.81
Buy and sell executed orders	+ 0.95
Opening and closing positions	+ 0.98

Table 4: Correlation coefficients computed for the different EUR/USD trading activity. Sampling period from January, 1st 2007 to January 31st 2009.

Statistic	Mean	Median	Minimum	Maximum	Std. Dev
Transaction numbers	1.45E+03	8.19E+02	0.00E+00	0 1.64E+04	1.81E+03
Transaction volumes	4.25E+07	1.41E+07	0.00E+00	9.01E+08	7.11E+07
Buy executed orders	7.22E+02	3.97E+02	0.00E+00	8.86E+03	9.20E+02
Sell executed orders	7.24E+02	4.04E+02	0.00E+00	9.07E+03	9.16E+02
Opening positions	8.20E+02	5.70E+02	0.00E+00	7.31E+03	8.58E+02
Closing positions	8.20E+02	5.61E+02	0.00E+00	7.34E+03	8.66E+02

Table 5: Descriptive statistics for the hourly aggregated (a) transaction numbers, (b) transaction volumes, (c) number of buy executed orders, (d) number of sell executed orders, (e) number of opening positions and (f) number of closing positions. Sampling period from January, 1st 2007 to January 31st 2009.

high positive linear relationship between (a) trade numbers and volumes, (b) numbers of buy and sell executed orders, and (c) numbers of opening and closing positions. The strength of the positive correlations is reflected in their correlation coefficients, with the results reported in Table 4.

In Table 5, we can see some of the statistical properties of the hourly aggregated (a) transactions numbers, (b) transactions volumes, (c) number of buy executed orders, (d) number of sell executed orders, (e) number of opening positions and (f) number of closing positions. The basic descriptive statistic properties reported in Table 5 are the mean, median, minimum, maximum and the standard deviation. From Table 5 we can see that there are extremely close similarities between the statistical properties for the number of buy and sell executed orders as well as for the number of opening and closing positions.

The volatility $vol_{L,t}$ is a measure of the variation of market price p_t at time

t over a period of length L . Given a price time series $\{p_t, t \geq 0\}$ and given a period of length L , the volatility is defined as:

$$vol_{L,t} = \frac{\sigma(p_t, \dots, p_{t-L+1})}{\frac{1}{L} \sum_{i=1}^L p_{t-i}} \quad (23)$$

where σ is the standard deviation.

In this study, we found that the EUR/USD mid-price volatility of the last 24 hours is correlated with the transaction numbers and volumes. The correlation coefficients computed for the EUR/USD mid-price volatility of the last 24 hours and the transaction numbers is +0.43 while for the transaction volumes it is + 0.28. We can thus hypothesise that an increase in the transaction numbers would lead to a rise in the price volatility, and vice versa. It is important to point out that this positive correlation does not mean that changes in the price volatility cause changes in the transaction numbers, and vice versa. To determine what causes changes in terms of the direction of the transaction numbers, an assessment of the overall market conditions must be carried out.

5 Conclusions

In this paper, we presented the study and analysis of the high frequency FX dataset provided by OANDA corporation with the aim being to establish the stylized facts of the FX market transactions data. We undertook an analysis of the FX market traders' collective behaviour which focused on: scaling laws, seasonality statistics and correlation behaviour.

Our contribution is two-fold. Firstly, using a HF dataset we have been able to confirm a number of stylized facts in the FX market that have been described in the literature. Secondly, our work goes beyond those and we have discovered four new scaling laws and six quantitative relations among them that apply to transactions data.

We have studied seasonality in the HF transactions dataset and have confirmed the existence of daily and weekly periodic patterns in the FX market transactions data. This adds to the studies of [1, 7] in the literature where they confirmed the existence of the double U-shape pattern of intraday transactions in FX markets [1, 7]. The intraday seasonality of transactions can

be explained by considering the behaviour of worldwide FX market centres whose business trading hours partially overlap. The intraweek seasonality can be explained by considering the daily patterns and the weekend effect. The intraday and intraweek seasonality show that a high level of transactions takes place when two or more FX market centre business hours overlap. Trading activity peaks during the opening business hours, declines during the lunch break, peaks again in the afternoon, then declines gradually during the closing business hours of an FX market centre.

We have investigated the correlation behaviour of high-frequency transactions data. This adds to other studies reported in the literature [1, 7]. Our study is different in that we have explored the correlation behaviour in different types of transaction in HF FX market data. It has been found that a high positive correlation exists between (a) transactions numbers and volumes, (b) numbers of buy and sell executed order, and (c) numbers of opening and closing positions. In addition, the price volatility of the last 24 hours is correlated with the activities measured in terms of the transactions numbers and volumes.

The second main contribution of our study is the discovery of four new scaling laws and six quantitative relationships among them (section 4.1.2). The statistical analysis of the scaling laws is based on the directional-change event approach which is event driven and the scaling laws apply to transaction data. Given a fixed threshold (Δx), the observed four independent new scaling laws state that the average (a) transaction numbers, (b) transactions volumes, numbers of (c) opening and (d) closing positions observed during a price move of size Δx is scale-invariant to the size of this threshold. The six quantitative relationships specify that, on average, an OS event contains roughly twice as many transaction numbers and volumes as a DC event. Furthermore, an OS event contains roughly twice as many numbers of opening and closing positions as a DC event. Additionally, an OS event contains roughly the same number of opening and closing positions, and the same features hold for a DC event. As far as we know this is the first such study that has studied and observed these relations. This adds to Glattfelder et al.'s work [20] where they uncovered 12 independent new scaling laws in foreign exchange price data.

The established stylized facts of the FX market transactions data could provide a foundation with regard to useful tools for forecasting price move-

ments and developing decision support systems. These stylized facts could allow us to make quantitative predictions regarding the behaviour of price movements and the dynamics of transactions flow in FX markets, in order to identify investment arbitrage opportunities. This could be done by predicting the next likely price peak or trough in the price time series. Moreover, the stylized facts can be fed into algorithmic trading strategies as a means of identifying patterns in the price and transaction time series such as to enable the strategies to adapt to changing market behaviour.

Stylized facts can be used as a benchmark to assess the ability of an artificial market to model the real market with a high degree of confidence. As part of our future work, we wish to understand the source of the stylized facts in the FX market, or how different elements of the market contribute to and affect their emergence. This will be done by developing an agent-based FX market which we will validate using the stylized facts reported in this paper.

Another avenue for future work is to establish more stylized facts of traders' behaviour in high-frequency FX markets, such as stylized facts of the traders' portfolios and their historical positions. This would involve establishing a foundation for modelling traders' behaviour from the microscopic analysis of the individual traders' transactions in the FX market. Another interesting line of future research is to identify the individual trader's adopted trading strategies through observing and analyzing the high-frequency dataset of OANDA individual traders' historical transactions. The aim would be to assess the trading strategies' performance by identifying and evaluating the strategies adopted. Accordingly, one could assess and classify the trading strategies that would lead to success or failure.

Acknowledgments

The authors would like to thank OANDA Corporation for providing the FX market high-frequency datasets.

References

- [1] M. Dacorogna, R. Genay, U. Mller, R. Olsen, O. Pictet, An introduction to high-frequency finance, Academic Press, San Diego, 2001.
- [2] D. Guillaume, M. Dacorogna, R. Dav, U. Mller, R. Olsen, O. Pictet, From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets, *Finance Stoch.* 1 (1997) 95–129.
- [3] R. Cont, Empirical properties of asset returns: stylized facts and statistical issues, *Quantitative Finance* 1 (2) (2001) 223–236.
- [4] B. LeBaron, A builder's guide to agent based financial markets, *Quantitative Finance* 1 (2) (2001) 254–261.
- [5] M. Evans, R. Lyons, Order flow and exchange rate dynamics, *Journal of Political Economy* 110 (2002) 170–180.
- [6] M. Evans, R. Lyons, Exchange rate fundamentals and order flow, Tech. Rep. 13151, NBER (2004).
- [7] T. Ito, Y. Hashimoto, Intraday seasonality in activities of the foreign exchange markets: Evidence from the electronic broking system, *Journal of the Japanese and International Economies*, Elsevier 20 (2006) 637–664.
- [8] J. Bouchaud, I. Giardina, M. Mzard, On a universal mechanism for long-range volatility correlations, *Quantitative Finance* 1 (2) (2001) 212–216.
- [9] Y. Chang, S. Taylor, Information arrivals and intraday exchange rate volatility, *Journal of International Financial Markets, Institutions and Money*, Elsevier 13 (2) (2003) 85–112.
- [10] R. Cont, Volatility clustering in financial markets: Empirical facts and agent based models, in: G. Kirman, A. & Teyssiere (Ed.), *Long memory in economics*, Springer, 2005.
- [11] F. Corsi, G. Zumbach, U. Mller, M. Dacorogna, Consistent high-precision volatility from high-frequency data, *Economic Notes - Review of Banking, Finance and Monetary Economics* 30 (2) (2001) 183 – 204.

- [12] G. Kim, H. Markowitz, Investment rules, margin, and market volatility, *Journal of Portfolio Management* 16 (1) (1989) 45–52.
- [13] S. Thurner, J. Farmer, J. Geanakoplos, Leverage causes fat tails and clustered volatility, *Quantitative Finance* (0908.1555).
- [14] T. Lux, The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions, *Journal of Economic Behavior & Organization* 33 (2) (1998) 143–165.
- [15] U. Mller, M. Dacorogna, R. Olsen, O. Pictet, M. Schwarz, C. Morgenegg, Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis, *J. Bank. Finance* 14 (1990) 1189–1208.
- [16] R. Mantegna, H. Stanley, Scaling behavior in the dynamics of an economic index, *Nature* 376 (1995) 46–49.
- [17] S. Galluccio, G. Caldarelli, M. Marsili, Y.-C. Zhang, Scaling in currency exchange, *Physica A* 245 (1997) 423– 436.
- [18] G. Balocchi, M. Dacorogna, C. Hopman, U. Mller, R. Olsen, The intraday multivariate structure of the euro futures markets, *Journal of Empirical Finance* 6 (1999) 479 – 513.
- [19] T. Di Matteo, T. Aste, M. Dacorogna, Long term memories of developed and emerging markets: using the scaling analysis to characterize their stage of development, *J. Bank. Finance* 29 (4) (2005) 827 – 851.
- [20] J. Glattfelder, A. Dupuis, R. Olsen, Patterns in high-frequency FX data: discovery of 12 empirical scaling laws, *Quantitative Finance* 11 (4) (2011) 599–614.
- [21] S. Ghashghaie, P. Talkner, W. Breymann, J. Peinke, Y. Dodge, Turbulent cascades in foreign exchange markets, *Nature* 381 (1996) 767 – 770.
- [22] D. Sornette, Fokker-planck equation of distributions of financial returns and power laws, *Physica A* 290 (1) (2000) 211–217.

- [23] X. Gabaix, P. Gopikrishnan, V. Plerou, H. Stanley, A theory of power-law distributions in financial market fluctuations, *Nature* 423 (6937) (2003) 267 – 270.
- [24] T. Di Matteo, Multi-scaling in finance, *Quantitative Finance* 7 (1) (2007) 21 – 36.
- [25] D. Berger, A. Chaboud, S. Chernenko, E. Howorka, J. Wright, Order flow and exchange rate dynamics in electronic brokerage system data, Tech. Rep. 830, Board of Governors of the Federal Reserve System, *International Finance* (2006).
- [26] J. Bouchaud, J. Farmer, F. Lillo, *Handbook of financial markets: dynamics and evolution*, Elsevier, 2009, Ch. How markets slowly digest changes in supply and demand, pp. 57–160.
- [27] M. Gould, M. Porter, S. Williams, M. McDonald, D. Fenn, S. Howison, *Limit order books*, *Quantitative Finance*.
- [28] K. Kim, S. Yoon, Y. Kim, Herd behaviors in the stock and foreign exchange markets, *Physica A: Statistical Mechanics and its Applications* 341 (2004) 526–532.
- [29] M. Aguirre, R. Said, Feedback trading in exchange-rate markets: evidence from within and across economic blocks, *Journal of Economics and Finance* 23 (1999) 1–14.
- [30] G. H. Bjnnes, D. Rime, Dealer behavior and trading systems in foreign exchange markets, *Journal of Financial Economics* 75 (2005) 571–605.
- [31] N. Laopodis, Feedback trading and autocorrelation interactions in the foreign exchange market: further evidence, *Economic Modelling* 22 (2005) 811–827.
- [32] J. Frankel, K. Froot, Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review* 80 (1990) 181–185.
- [33] J. Frankel, K. Froot, Exchange rate forecasting techniques, survey: data, and implications for the foreign exchange market, Tech. Rep. WP 90, International Monetary Fund, Washington, DC (1990).

- [34] T. Ito, Foreign exchange rate expectations: micro survey data, *American Economic Review* 3 (1990) 434–449.
- [35] R. MacDonald, I. Marsh, Currency forecasters are heterogeneous: confirmation and consequences, *Journal of International Money and Finance* 15 (1996) 665–685.
- [36] T. Oberlechner, Evaluation of currencies in the foreign exchange market: attitudes and expectations of foreign exchange traders, *Zeitschrift fuer Sozialpsychologie* 32 (3) (2001) 180–188.
- [37] L. Menkhoff, R. Rebitzky, M. Schroder, Heterogeneity in exchange rate expectations: evidence on the chartist-fundamentalist approach, *Journal of Economic Behavior and Organization* 70 (2009) 241–252.
- [38] S. Masry, M. ALoud, A. Dupuis, R. Olsen, E. Tsang, High frequency FOREX market transaction data handling, in: 4th CSDA International Conference on Computational and Financial Econometrics, London, UK, 2010.
- [39] Y. Bingcheng, E. Zivot, Analysis of high-frequency financial data with S-PLUS, Tech. Rep. UWEC-2005-03, University of Washington, Department of Economics (2003).
- [40] G. Brownlees, C. and Gallo, Financial econometric analysis at ultra-high frequency: data handling concerns, *Computational Statistics & Data Analysis* 51 (2006) 2232–2245.
- [41] M. Blume, M. Goldstein, Quotes, order flow, and price discovery, *Journal of Finance* 52 (1) (1997) 221–244.
- [42] R. Oomen, Properties of realized variance under alternative sampling schemes, *Journal of Business & Economic Statistics* 24 (2006) 219–237.
- [43] E. Tsang, Directional changes, definitions, Tech. Rep. 050-10, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, UK (2010).

- [44] M. Aloud, E. Tsang, R. Olsen, A. Dupuis, A directional-change events approach for studying financial time series, *Economics Papers* No 2011-28.
- [45] S. Fong, Y. Si, J. Tai, Trend following algorithms in automated derivatives market trading, *Expert Systems with Applications* 39 (13) (2012) 11378–11390.
- [46] M. Covel, *Trend following: How great traders make millions in up or down markets*, Financial Times Prentice Hall, 2004.
- [47] A. Szakmarya, Q. Shenb, S. Sharmac, Trend-following trading strategies in commodity futures: A re-examination, *Journal of Banking & Finance* 34 (2) (2010) 409–426.
- [48] A. Sansone, G. Garofalo, Asset price dynamics in a financial market with heterogeneous trading strategies and time delays, *Physica A* 382 (2007) 247–257.
- [49] L. Harris, Transaction data tests of the mixture of distributions hypothesis, *Journal of Financial and Quantitative Analysis* 22 (2) (1987) 127–141.

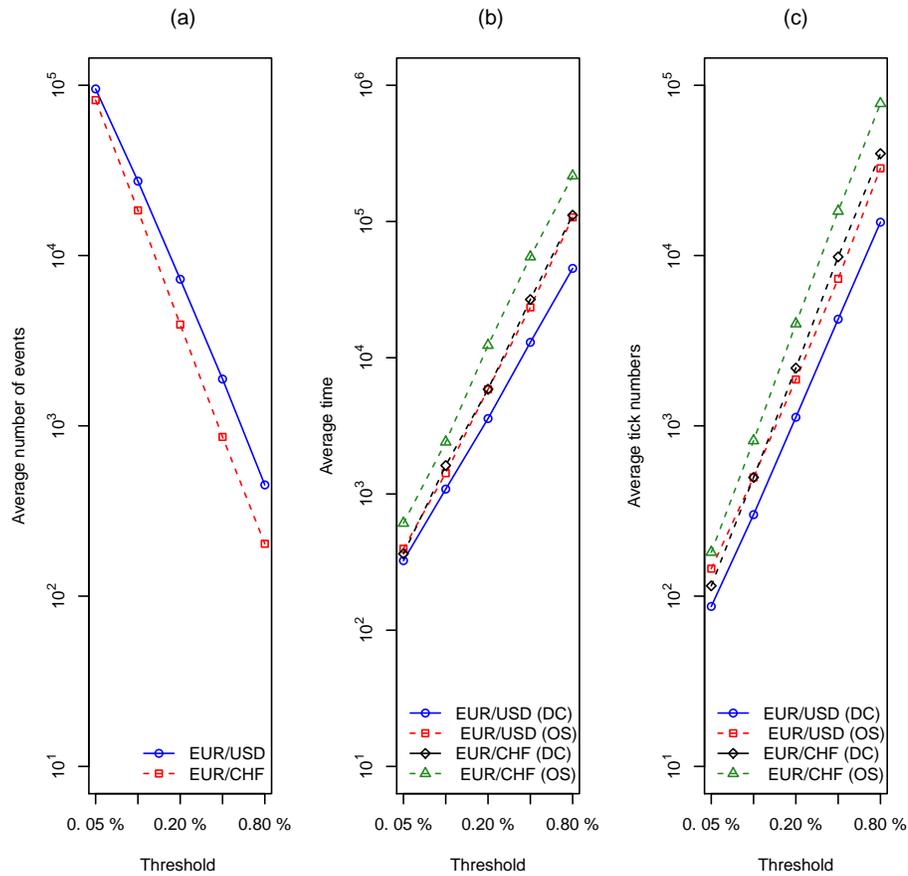


Figure 1: Scaling laws (a), (b) and (c) are plotted where the x-axis shows the price moves thresholds of the EUR/USD and EUR/CHF observations and the y-axis the (a) average number of events, the (b) average time (in seconds) and the (c) average tick numbers.

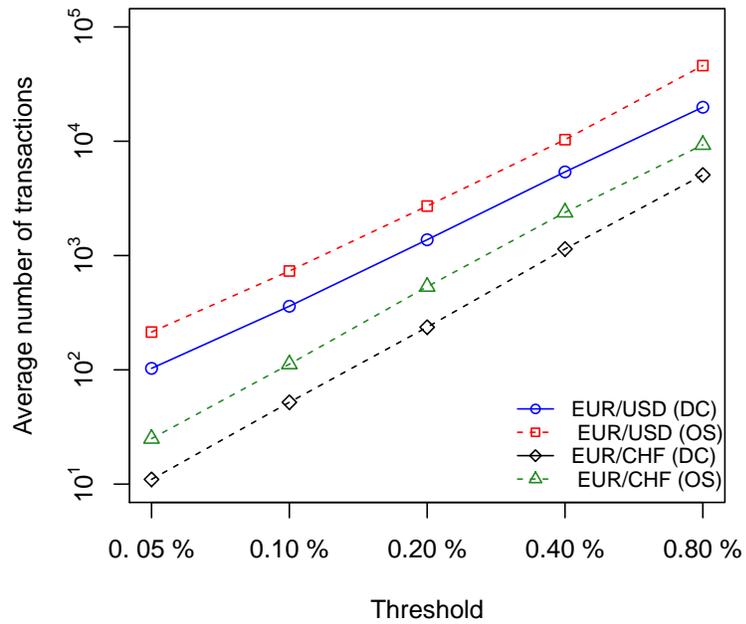


Figure 2: Average EUR/USD and EUR/CHF transaction (trade) numbers during DC and OS events for selected threshold values. The x-axis shows the price moves thresholds of the EUR/USD and EUR/CHF observations and the y-axis the average transaction numbers.

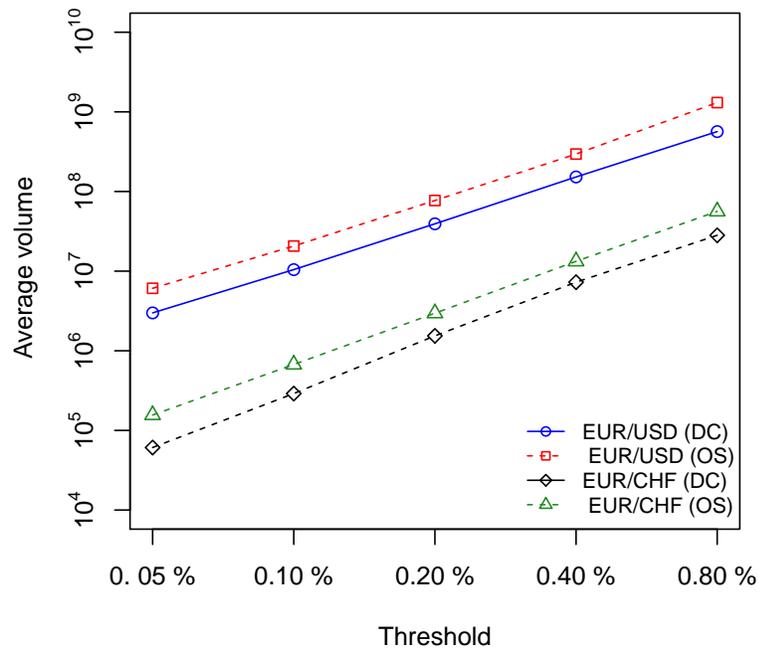


Figure 3: Average EUR/USD and EUR/CHF transaction volumes during DC and OS events for selected threshold values. The x-axis shows the price moves thresholds of the EUR/USD and EUR/CHF observations and the y-axis the average transaction (trade) volumes.

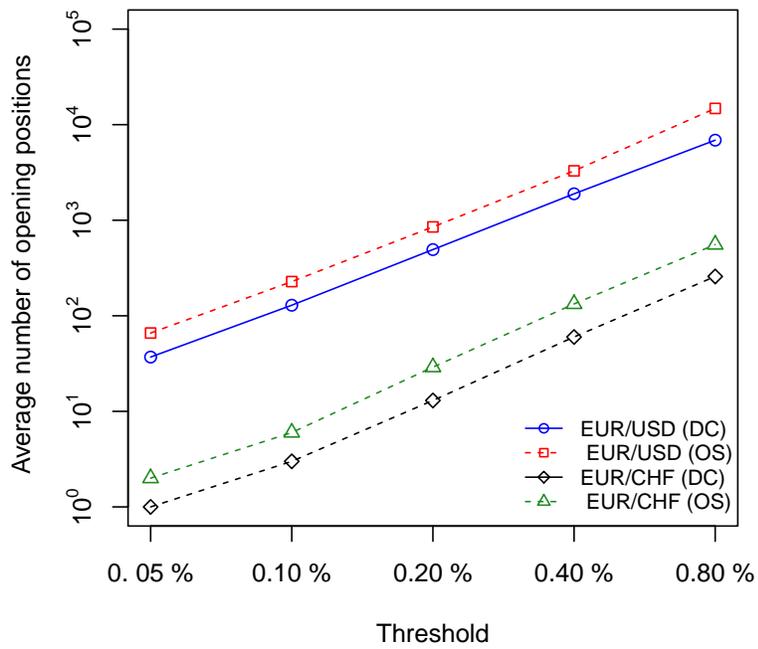


Figure 4: Average number of opening EUR/USD and EUR/CHF positions observed during DC and OS events for selected threshold values. The x-axis shows the price moves thresholds of the EUR/USD and EUR/CHF observations and the y-axis the average number of opening positions.

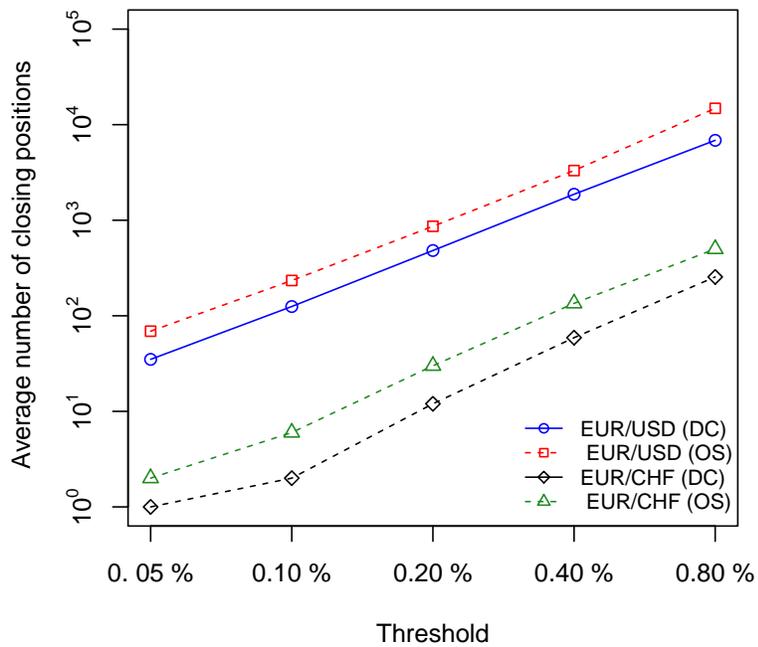


Figure 5: Average number of closing EUR/USD and EUR/CHF positions observed during DC and OS events for selected threshold values. The x-axis shows the price moves thresholds of the EUR/USD and EUR/CHF observations and the y-axis the average number of closing positions.

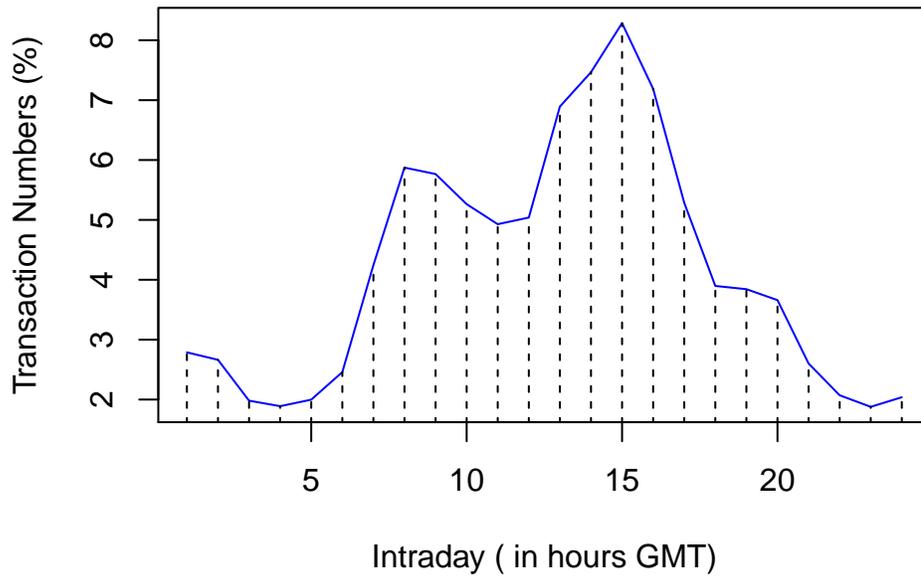


Figure 6: Hourly intraday seasonality of the transaction numbers showing the average transaction numbers in each hour of the day. The sampling interval is one hour ($\Delta t = 1$). The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009. The transactions are made in terms of 48 different currency pairs. The time scale is GMT.

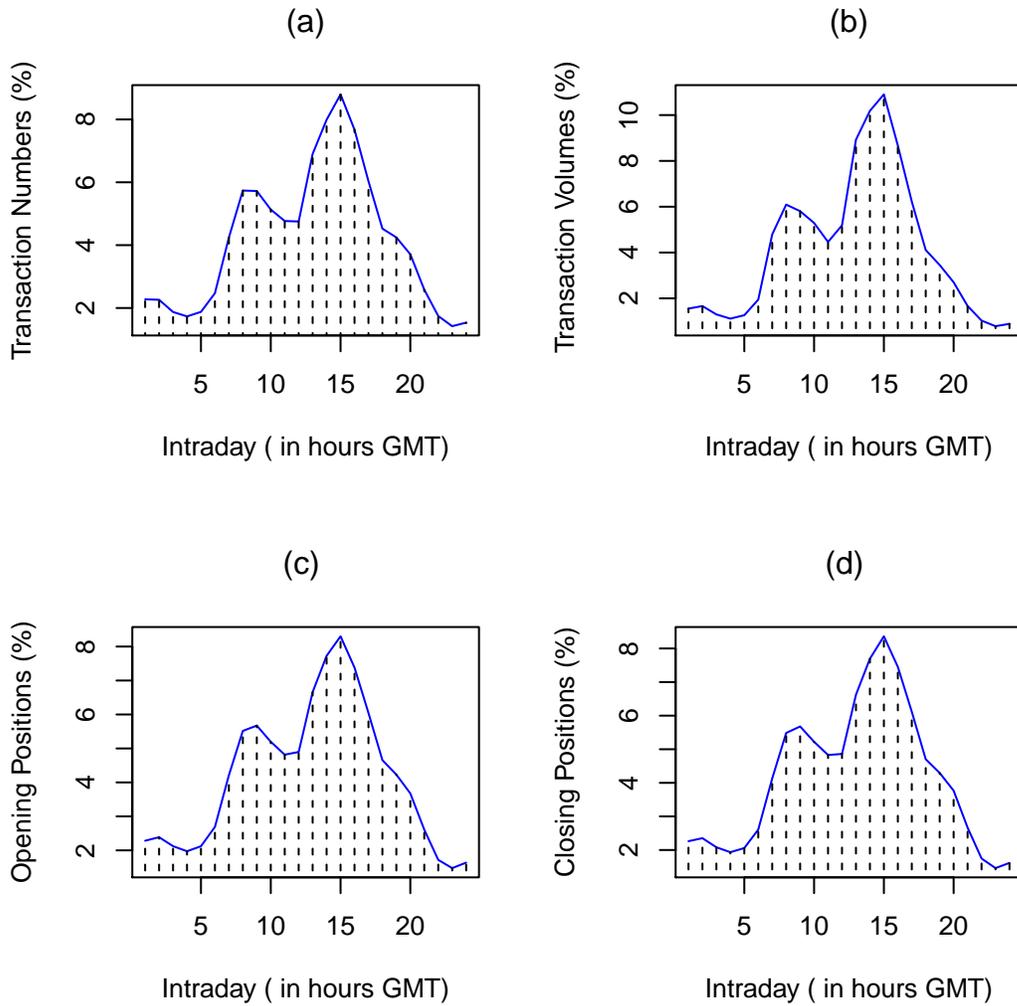


Figure 7: Hourly intraday seasonality of EUR/USD (a) transaction numbers, (b) transactions volumes, (c) number of opening positions and (d) number of closing positions. The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009. The time scale is GMT. All the four plots show similar double U-shapes.

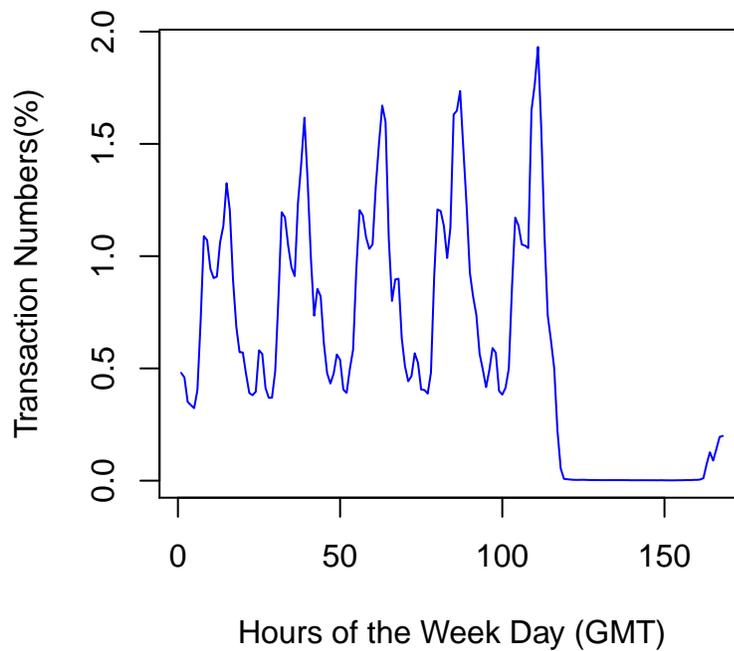


Figure 8: Hourly intraweek seasonality of the transaction numbers showing the average number of transactions in each hour of the weekday. The sampling interval is one hour ($\Delta t = 1$). The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009. The transactions are made in terms of 48 different currency pairs. The time scale is GMT.

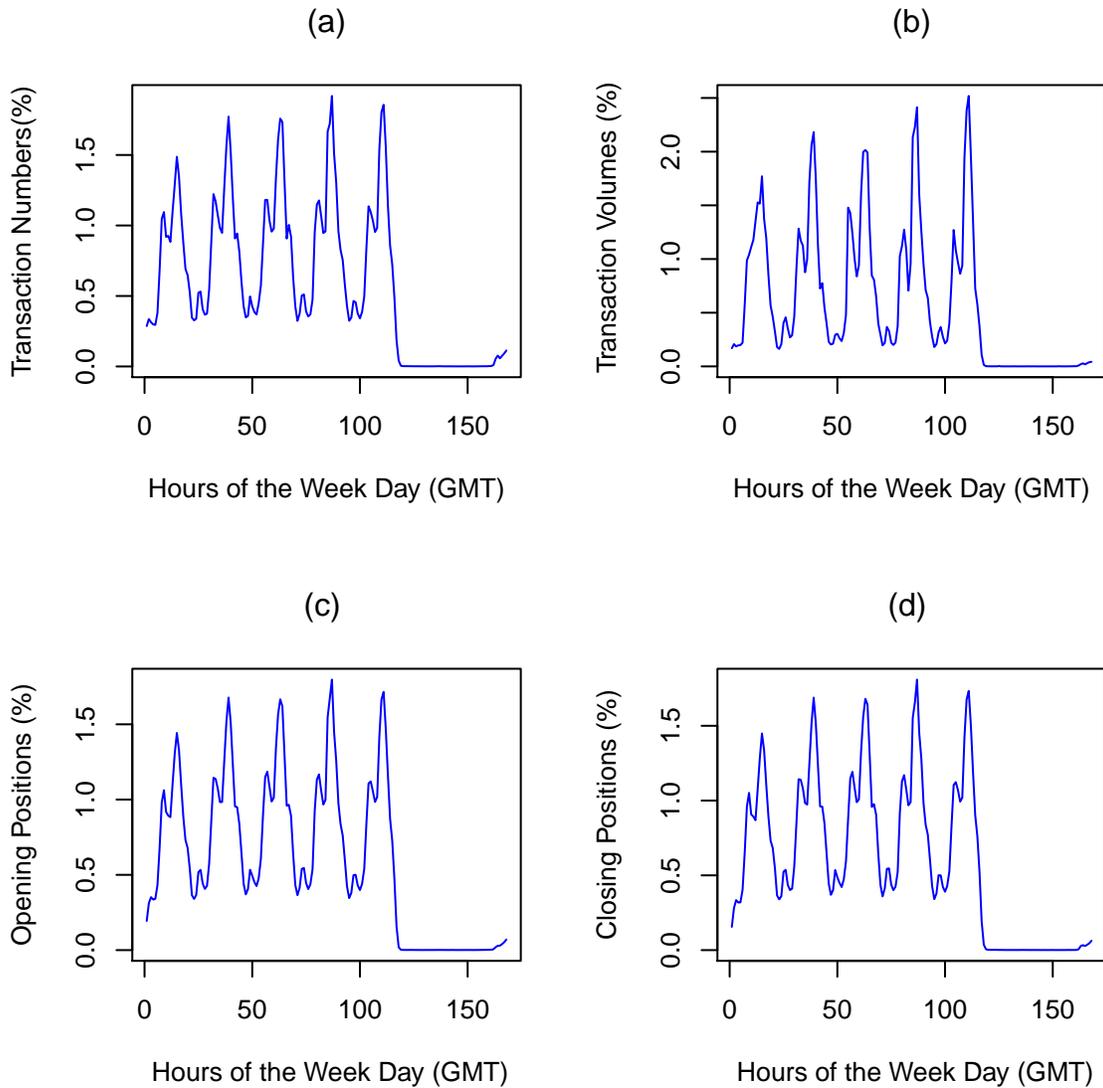


Figure 9: Hourly intraweek seasonality of EUR/USD (a) transactions numbers, (b) transactions volumes, (c) number of opening positions and (d) number of closing positions. The sampling period covers 2.25 years from 1st January 2007 to 5th March 2009. The time scale is GMT.

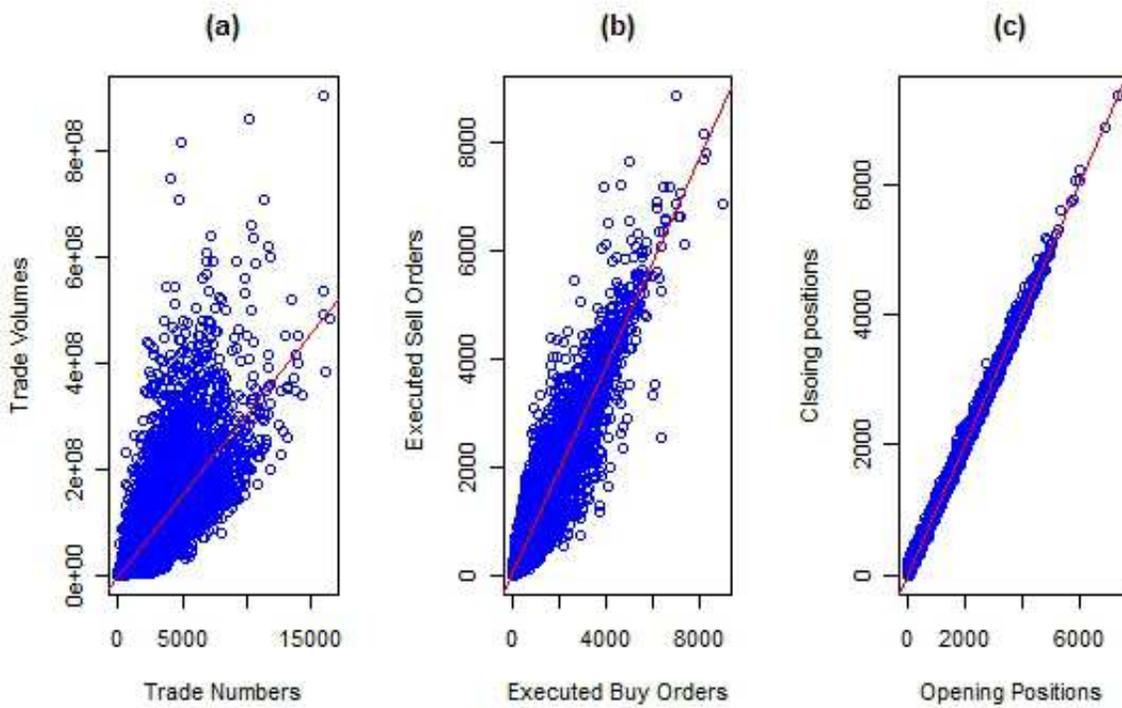


Figure 10: Correlation of EUR/USD (a) transaction (trade) numbers and volumes, (b) numbers of buy and sell executed orders, and (c) numbers of opening and closing positions: a sampling interval of $\Delta t = 1$ hour is chosen. The sampling period from January, 1st 2007 to January 31st 2009.