

Modelling the High-Frequency FX Market: an Agent-Based Approach

Monira Aloud^{1,3}, Maria Fasli^{2,3}, Edward Tsang^{2,3}, Alexander Dupuis^{3,4} and Richard Olsen^{3,4}

September 4, 2016

¹ Department of Management Information Systems, College of Business Administration, King Saud University, KSA

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

³ Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, UK

⁴ Olsen Ltd., Seefeldstrasse 233, 8008 Zurich, Switzerland

Email: mealoud@ksu.edu.sa, mfasli@essex.ac.uk, edward@essex.ac.uk, alex@olsen.ch and richardo@olsen.ch

Abstract

The development of computational-intelligence based strategies for electronic markets has been the focus of intense research. In order to be able to design efficient and effective automated trading strategies, one first needs to understand the workings of the market, the strategies that traders use and their interactions as well as the patterns emerging as a result of these interactions. In this paper, we develop an agent-based model of the FX market which is the market for the buying and selling of currencies. Our agent-based model of the FX market (ABFXM) comprises heterogeneous trading agents which employ a strategy that identifies and responds to periodic patterns in the price time series. We use the ABFXM to undertake a systematic exploration of its constituent elements and their impact on the stylized facts (statistical patterns) of transactions data. This enables us to identify a set of sufficient conditions which result in the emergence of the stylized facts similarly to the real market data, and formulate a model which closely approximates the stylized facts. We use a unique high frequency dataset of historical transactions data which enables us to run multiple simulation runs and validate our approach and draw comparisons and conclusions for each market setting.

Keywords: agent-based modelling, agent-based simulation, electronic markets, FX markets, stylized facts.

1 Introduction

The massive expansion of the Internet and developments of related information and communication technologies have had a considerable impact on almost all aspects of human life, as well as on business and industry sectors. Financial and other markets have seen an increase in the deployment of automated software, so-called algorithmic traders (Booth et al., 2014), that participate in the markets and trade alongside human traders. In the not so distant future, marketplaces can be envisaged as being populated by autonomous intelligent entities – software/trading agents – which represent their owners (be it individuals, corporations or other organisations) and conduct business on their behalf. Unlike “traditional” software, agents are personalised, semi-autonomous and continuously running entities able to adapt to conditions in a dynamic environment (Fasli, 2007).

Although developing such software agents for clearly defined tasks in static environments can be straightforward, electronic markets are complex and volatile environments and developing agents with flexible strategies able to cope with uncertainty and able to adapt is non-trivial. In such dynamic environments, populated by thousands of entities, successful performance of a trading agent does not only depend on its strategy, but on the strategy of the other agents and entities as well.

To be able to design efficient and effective trading strategies, one first needs to understand the workings of the market, the strategies that traders use and their interactions as well as the patterns emerging as a result of these interactions. But this is often hampered by two difficulties (Fasli and Michalakopoulos, 2008). Although “the rules of the game” (market rules) according to which traders trade are known, one may not be able to obtain sufficiently large and detailed datasets from the real market to analyse in order to understand market behaviour at the macro level. Even if one obtains real data from a market, real traders are in general reluctant to reveal the strategies they employ when operating in markets for fear of losing their competitive advantage. Hence, it is very difficult to study traders’ behaviour at the micro level which inherently hinders understanding the link between micro and macro behaviour.

The aim of our work is two-fold. On the one hand, we want to understand market behaviour in order to develop flexible and effective trading strategies based on artificial intelligence techniques. On the other, by understanding market behaviour and patterns as they emerge from the interaction of traders, we wish to understand the emergence of phenomena such as **crashes** and the conditions and underlying behaviours that may lead to these events and phenomena. This is also invaluable from a policy maker or market designer perspective: understanding conditions that may lead to phenomena such as disruptions in the market could help us develop decision support systems acting as early warning systems that could potentially enable policy makers/market designers put in place corrective action to avoid catastrophic consequences. Our specific interest is in Financial Markets and the Foreign Exchange (FX) market in particular, which is the market for the buying and selling of currencies.

In studying financial markets and understanding their behaviour, the first key step is establishing their statistical properties, otherwise known as *stylized facts* (Dacorogna et al., 2001) which are in essence the patterns that emerge in the market at the macro level as a result of the interactions of traders and their behaviours at the micro level. The second key step is understanding the origins of the stylized facts, i.e. what elements may lead to or affect their emergence or what are the underlying individual behaviours that in essence result in these patterns. As mentioned above, even if one has data from the real market and is able to identify stylized facts, obtaining information about individual traders’ behaviour and underlying strategies is far more difficult. To this end, typically one formulates hypotheses as to the potential origins of the stylized facts.

Ideally, the next step then should be to validate these hypotheses. Exploring and understanding the origins of stylized facts in a market setting is intricate as one cannot usually control various aspects of the real market. Hence, the only conceivable and feasible means of studying the origin and emergence of stylized facts is via modelling the market and undertaking an exploratory study. Due to the complexity of real markets, especially high frequency markets such as the Foreign Exchange market, traditional analytical models present difficulties in analysing financial market behaviour (Martinez-Jaramillo and Tsang, 2009). Agent-based models of markets (ABMs) have emerged as a powerful tool in the exploration and understanding of markets and trading behaviour. ABMs are computer simulations that comprise different elements such as a mechanism (trading rules), traded assets and trading agents. ABMs have been used in a range of studies from macroeconomic models and payment card markets to advert markets and retail electric energy markets. Hence agent-based models can be used to emulate a market with the aim being to develop a model based on the hypotheses formulated regarding the potential origins of the stylized facts and the individual traders’ behaviours that closely approximates the real

market. Having a model of an artificial market that closely approximates the stylized facts and patterns as exemplified in the real market data serves as a validation step for the hypotheses. This provides a higher degree of confidence in the understanding of the market workings and though one may not be able to identify all necessary conditions for the emergence of the stylized facts, the approximation would provide evidence that a set of sufficient conditions has been identified that results in the emergence of the stylized facts.

Once you have in place a model that approximates the real market workings and stylized facts, the subsequent step is to develop and experiment with alternative strategies and explore *what-if* scenarios to obtain a deeper understanding of both the market workings as well as how the overall market behaviour would evolve and adapt to new conditions. We can also use this understanding to devise efficient and effective trading strategies and test them under various conditions in the artificial market and in comparison to other strategies.

In this paper, we develop and demonstrate an agent-based model of the Foreign Exchange (FX) market. Our study has been supported through the provision of a High Frequency dataset by OANDA corporation which contains over 147 million transactions over a period of more than two years by over 45,000 trading accounts (anonymised). This has enabled us to study and identify the stylized facts of the FX market (Aloud et al., 2013) which we use in this paper in order to arrive at a model of the artificial market that approximates the stylized facts of the real market data.

We develop a simple and flexible agent-based FX market (ABFXM) which is populated by a number of heterogeneous interacting agents. The agents employ the zero-intelligence directional-change event trading strategy (ZI-DCT0) and are able to identify and respond to periodic patterns in the price time series. One of our fundamental hypothesis is that agents in such a market are heterogeneous traders **and hence in our ABM, heterogeneity is modelled in the agents' initial wealth, their trading strategies, their profit objectives, their risk appetites and their trading windows which represent geographical distribution.** Furthermore, the ABFXM has a number of other distinguishing features including asynchronous trading windows and an approach for determining an agent's order size and limit orders. Using the stylized facts which were identified from the real HF FX dataset (Aloud et al., 2013), we evaluate the transactions produced from the ABFXM in resembling the transactions of the real FX market. By relating the stylized facts of actual FX market traders' behaviour to the agent-based models of traders, we may find a pathway to identify a set of elements sufficient for modelling the transactions in the FX market. We vary the constituent elements of the ABFXM in a systematic way, in order to explore their effect on the emergence of the stylized facts in the ensuing market behaviour.

The organization of the paper is as follows. Section 2 reviews the literature and discusses the issues that have not been sufficiently addressed yet in previous work in ABMs. Section 3 provides an overview of the FX market structure. The ABFXM is described in section 4. Section 5 presents the validation of the ABFXM in reproducing the stylized facts of the transactions in the FX market. In section 6, we present a series of controlled experiments which examine the implications of modifying elements of the ABFXM and we discuss the results. The paper ends with the conclusions and avenues for future work.

2 Background

Developing and testing strategies for software agents suitable for complex markets is fraught with difficulties. Firstly, to be able to develop effective and efficient strategies, one needs to have an excellent understanding of the environment in which the agent will be embedded. Such an understanding is essential for developing – as much as possible – an agile and adaptive strategy that will be able to cope with a highly dynamic and uncertain environment populated by other entities whose behaviour cannot be predicted in advance. Secondly, the cost of testing strategies in real markets can be prohibitively high, if the strategy is defective in any way or even if it has simple logical errors, hence realistically this is out of the question. Thirdly, even if one was willing and able to test one's strategy at a particular moment in time and achieve good performance, the problem is that measuring an agent's performance at just a snapshot in time may not represent the true performance of a strategy and in order to have confidence in the strategy's effectiveness, this should be tested for long periods of time; this will in essence pit the strategy against multiple other opponents and their strategies. However, inevitably this may lead to considerable **losses**.

To tackle this problem, the use of simulations has been advocated for the purposes of modelling complex environments and for testing new operating procedures – or strategies as in our case – especially where the cost of testing them within a real environment can be prohibitively high (Verma and Kenneth, 2010; Fasli and

Michalakopoulos, 2008). Computer simulation is an appropriate tool for process analysis especially if (Verma and Kenneth, 2010):

- the system dynamics are extremely complex;
- the system has one or more interdependent random variables;
- the objective is to observe system behaviour over a period of time.

Financial markets are complex systems which produce immense amounts of data. Due to the current large volume and complexity of financial market HF data, compared to earlier times, traditional analytical models present challenges in analysing such data and their use may not be able to yield satisfactory results and useful insights. Although there are some microstructure models which have studied the behaviour of the FX market, there has been no systematic study of the transactions properties using a HF transactions dataset and an agent-based model. In the financial market literature, models of the FX markets have focused on the price behaviour and acquired either large samples of low frequency datasets (Evans and Lyons, 2004), or small samples of HF datasets (Berger et al., 2006), neither of which are at the account level. Hence although these works provide an analysis of price behaviour, they offer no insight on the underlying behaviour of the traders and they were constructed based on low frequency data. The market dynamics cannot be explained only by macroeconomic data and analytical models. To understand the FX market workings, a systematic exploration of the properties of the transactions needs to be undertaken and this requires a HF dataset of individual traders' historical transactions over a substantially long period of time. Exploring and studying these micro data of the FX market will provide an insight into the collective macro-behaviour of FX trading activity and can help explain market phenomena.

Agent-based models of markets (ABMs) have emerged as an alternative and powerful modelling tool for studying markets and testing strategies. The basic units of an ABM are agents that can interact dynamically with each other, react to price changes and events in the market and may adapt to their environment and learn from the information they obtain. (LeBaron, 2001) offers a good overview of some of the most prominent ABMs. The area of ABMs has witnessed a rapid increase in the number of works produced since the introduction of the influential Santa Fe Artificial Stock Market (SF ASM) (Arthur et al., 1997). Some of the subsequent ABMs differ from the original SF ASM in terms of the type of trading agents involved (Chen and Yeh, 2001; Martinez-Jaramillo and Tsang, 2009); the market mechanism (Levy and Solomon, 1996; Lux, 1998); and the number of traded assets (Sunder, 2004). ABMs have also been used as the underlying platform for testing strategies in complex settings such as dynamic supply chain management, advert markets and electric energy markets (TAC, 2016). Our particular interests lie in financial markets and specifically the Foreign Exchange market.

One of the key elements in studying financial markets and understanding their behaviour is establishing their statistical properties, otherwise known as *stylized facts* (Dacorogna et al., 2001), and exploring their origins, i.e. what elements may lead to or affect their emergence. As stylized facts are established from the statistical analysis of real market data, this very much depends on the availability of such data. Stylized facts can also serve as a benchmark for validating that the modelled market exhibits the same properties as the real one (LeBaron, 2001).

Despite the existence of previous and contemporary studies on modelling financial markets using an ABM approach, a number of issues have not been well-studied or properly addressed as yet. ABMs have been typically aimed at reproducing some of the stylized facts of financial markets. However, many of these studies pay little attention to identifying and understanding the origins of the market stylized facts. For instance, the authors of the SF ASM do not provide a quantitative description of the reproduced stylized facts (Arthur et al., 1997). Examples of ABMs that attempt to explain the origin of stylized facts of the price returns include (Daniel, 2006; Alfi et al., 2009a; Martinez-Jaramillo and Tsang, 2009). The work in (Alfi et al., 2009a) presents an ABM containing the simplest possible elements to define and understand the origin of the stylized facts of price returns. In (Martinez-Jaramillo and Tsang, 2009), the authors model a simple market mechanism of a stock market which operated with sophisticated genetic programming based agents. Their work identifies the minimal set of essential conditions under which the stylized facts of endogenously generated prices resemble those of real market prices. (Daniel, 2006) focuses on the intraday stylized facts of order flow and price returns in the stock market and studies the origin of the stylized facts using a double auction market mechanism

operated with zero-intelligence agents. Our work is different from the works in (Daniel, 2006; Alfi et al., 2009a; Martinez-Jaramillo and Tsang, 2009), in that we used an agent-based approach to model trading behaviour in the high-frequency FX market and investigate the emergence of the stylized facts of transactions data in a systematic way.

Although ABMs offer a number of advantages in modelling markets, the use of such models has been criticized from a number of viewpoints and in particular with regard to the complexity involved in their setup and the problem of calibration (LeBaron, 2001). Simplicity is essential in designing ABMs (LeBaron, 2006; Hommes, 2006; Samanidou et al., 2007) as a complex design may impede the study, understanding and interpretation of the origins of the stylized facts (Hommes, 2006). Although it is often tempting to develop ABMs with intricate set ups and numerous composite parameters, such often unwarranted, complexity makes it difficult to identify which elements of the market are accountable for the emergence of the stylized facts and whether all of the elements are equally important for reproducing them (LeBaron, 2006). Identifying and understanding the origins of stylized facts is important for understanding the dynamics of market behaviour, so that we can develop computational intelligence approaches and strategies for software trading agents. Furthermore, understanding the origins of stylized facts helps shed light in the characteristics of a market from a policy-maker perspective; this can be useful in designing new policies and market structures and understanding their potential impact on the market and the participants' behaviour.

The contribution of this paper is fivefold. Firstly, we have developed an agent-based model of the FX market (ABFXM). To the best of our knowledge, there are no other works in the literature that have modelled the FX market using an agent-based approach. Our second contribution is in the construction of the agent-based model itself which comprises heterogeneous trading agents operating in an asynchronous manner. The third contribution is in the use of a strategy based on the directional change event approach (Aloud et al., 2011), the ZI-DCT0, which is able to identify and respond to periodic patterns in the price time series. The fourth contribution is in the use of the ABFXM to explore and study the origins of the stylized facts in the FX market. We have undertaken a comprehensive exploration of the constituent elements of the ABFXM and their impact on the stylized facts of transactions data and have identified a set of sufficient conditions which result in the emergence of the stylized facts similarly to the real market data. The fifth contribution is in the use of a unique high frequency dataset of historical transactions to undertake this study. We have been able to systematically study the effect of each of the constituent elements of the ABFXM on the stylized facts through anchoring the simulation results by using a historical high-frequency prices dataset. The use of this dataset enables us to run multiple simulation runs and draw comparisons and conclusions for each market setting against the real data but also compare different market settings against each other. As far as we know, this has not been done before.

3 Overview of the FX Market

The Foreign Exchange market, also referred to as the Forex or FX market (Dacorogna et al., 2001), is where the buying and selling of foreign currencies takes place. In contrast to other financial markets that imply an exchange of a valuable asset for a sum of money, the FX market implies an investment of one currency for another currency, hence in the FX market, each transaction entails a pair of currencies, e.g. EUR/USD. The FX market is a considerable market whose transactions averaged \$5.3 trillion per day in April 2013. To provide a measure of the FX market's size, Germany's GDP (4th biggest economy in the world) was \$3.7 trillion in 2013 while the UK's (6th economy in the world) was \$2.67 trillion for the same year.

In contrast to stock markets, the FX market is decentralized and this means that the FX market is not associated with specific physical locations. Furthermore, the FX market operates 24 hours a day, 7 days a week, with participants being from all over the world. Despite the fact that the FX market operates continuously on a 24-hour basis, the trading activity is not homogenous in time, which leads to difficulties in analyzing the data (Dacorogna et al., 2001). The peak liquidity occurs when the business trading hours of multiple time zones overlap and this is a well-known pattern (Dacorogna et al., 2001). It is essential to understand the correlation between market activity and liquidity over different periods of the day.

The rapid advances in communications technology, particularly in web-related technology, in the last two decades have led to huge changes in the FX market in that a number of independent brokers have developed internet-based trading platforms – the so-called market-maker market. A market-maker is an institution which provides liquidity for a number of currency pairs, and quotes both a buy and a sell price on its platform in

an attempt to make a profit through the bid/ask spread. In essence, a market-maker is an intermediary in the market and facilitates transactions. In an FX trading environment, the market-maker is equipped to buy from and sell to its investors and other market-makers at those quoted bid and ask prices. In other words, the market-maker takes the opposite side of a trade. The market-maker's role is essential in maintaining the liquidity with regard to a given security that is held in their inventory. A market-maker adjusts the price based on the demand and supply in the market, on their inventory holdings, on news and prices from other market-makers, as well as other factors. The role of the electronic market-maker has become more prominent in the last decade and electronic market-makers play an increasingly important role in the growth of the financial markets and the FX market in particular.

4 Agent-Based FX Market

The agent-based FX market (ABFXM) developed is populated by a number of agents who participate in the market in terms of buying and selling currencies. The traded asset is a currency pair. The ABFXM uses historical prices for the EUR/USD currency pair which are fed into the market through the market-maker component of the simulation and the agents react based on this information. At any time t in the market, an agent a can hold two different types of assets: a risk free asset in the form of cash and a risky asset in the form of a currency. The decision behind having just two different types of assets was based primarily on the complexity of the analysis of markets in which more than two assets are available for trading.

Prior to the start of the simulation, based on a continuous uniform distribution, each agent will be assigned the following elements: home currency, leverage, business trading hours corresponding to a market geographical location and profit objective and risk appetite thresholds. Each individual agent will be endowed with a different amount of cash (wealth) expressed in the agent's home currency. The agent design and market clearing mechanisms are described in more detail in the subsequent sections.

4.1 Agent design

In designing trading agents, there are two equally important but competing principles that we need to keep in mind. On the one hand, simplicity is key as a complex agent and in more general market design may impede the study, understanding and interpretation of the origins of the stylized facts (Hommes, 2006) and the dynamics of the market.

On the other hand, heterogeneity is inherent in real markets and they are populated by agents that have different preferences, demands, aspirations, trading strategies and trading hours. Other works have modelled heterogeneity in a limited way and through a small number of elements. In our work, we have attempted to satisfy both of these principles through a well-defined parameter space and clear strategy. In our ABFXM, heterogeneity is modelled through: (a) wealth variations among the agents (**Section 4.1.1**); (b) different behavioural groups of traders (**contrarian and trend-following agents; Section 4.1.4**); (c) agents who have different profit objectives (**Section 4.1.3**); (d) agents who have different risk appetites that form long and short term investors (**Section 4.1.3**); and finally (e) agents who use different trading time windows (**Section 4.1.7**). All of these aspects build heterogeneity into the trading model, leading to the characterization of different trading activities. The following sections describe the various elements of the agents' design and their by default setting.

4.1.1 Initial wealth distribution

The wealth represents the amount of cash that an agent holds and it is expressed in the home currency of the agent's account. Many ABMs imply that all agents are endowed with the same amount of cash prior to the launch of the simulation, as has been done in (Levy et al., 1994; Martinez-Jaramillo and Tsang, 2009). However, this is an unrealistic assumption: real markets involve heterogeneity in terms of the wealth distribution of the traders and we would argue that endowing all traders with the same wealth and hence giving them the same market power is not representative of what happens in reality. We would therefore like to put forward the case for differentiating the initial wealth allocation to traders which is more in line with real life.

Over a century ago, Pareto found through empirical observation, that the income distribution in numerous countries follows a power-law distribution (Pierre, 1965). In the follow-up to Pareto's study, a number of

empirical studies have supported this thesis by providing evidence from empirical data (Steindl, 1965). We quantify the income distribution in financial markets by a power-law distribution to endow each agent with a different amount of cash. The range of the distribution is a critical element. Estimating the values of the range and the scaling exponent for the power-law distribution are very important due to their role in characterising the distribution. Hence, estimating the range and the scaling exponent needs to be done carefully. For instance, the larger the scaling exponent, the steeper the curve, consequently implying more small scales than large ones. **Using a systematic approach, a sensitivity analysis has been conducted to examine the impact of different values for the scaling law component value and for the range of agents' wealth in the exhibition of the FX market transactions' stylized facts.**

The power-law distribution is employed in the simulation by converting a random number that is uniformly distributed into a number that is distributed along the power-law distribution. The power-law and the random variable are related via:

$$\int_{l_0}^{l_1} k.l^{-\alpha} = \int_0^z 1 dz \quad (1)$$

where l_0 and l_1 are the lower and upper bounds respectively from which the power law distribution holds. k is a constant value defined by:

$$k = \frac{-\alpha + 1}{l_1^{(-\alpha+1)} - l_0^{(-\alpha+1)}} \quad (2)$$

α is the scaling exponent for a power-law distribution (where $\alpha = 1.2$) and z is a random number drawn from a **uniform distribution** using a specified range of values [10^3 , 10^9]. Combining these two equations, the power law has the following form:

$$l_z = \left(\frac{z(-\alpha + 1)}{k} + l_0^{(-\alpha+1)} \right)^{(-\alpha+1)} \quad (3)$$

We would argue that the assumption that agents possess the same wealth at the beginning of the market as has been done in (Daniel, 2006; Martinez-Jaramillo and Tsang, 2009; Levy et al., 1994; LeBaron, 2002; Kim and Markowitz, 1989), may require reconsideration in this area. As our work demonstrates, assuming a uniform distribution of wealth amongst traders (at least in this specific market), does not seem to support the emergence of the stylized facts of the FX market (Section 6.4).

4.1.2 Margin trading

Trading on margins enables a trader to place an order when its size is larger than its cash holding. This involves borrowing a fraction of the cash required from the market-maker executing the transaction and depends on the leverage ratio used by the trader. A leverage ratio of $v:1$ requires a minimum margin $m\%$ ($\frac{1}{v}$), hence, if a trader decides to open a new position, it must have at least $m\%$ of the position's size available as a margin. Trading on the margin may possibly realize high profits, but this in turn involves high risk. Margin trading makes use of a course of action referred to as a margin call (Section 4.2).

4.1.3 Profit objective and risk appetite

Using the HF individual traders' transactions dataset described in section 5.1, we map the traders' profit objectives and risk appetites into the agent-based models of traders. The mapping was accomplished by tracking each trader's transactions in each of their traded currency pairs. This is done in terms of tracking the number of currency units bought and sold sequentially by the transactions' execution time. Consequently, we trace a trader's different types of orders with regard to a given currency pair from the opening to the closing of a position. This allows us to define the trader's profit and loss thresholds by which the trader decides to open and close their positions. The threshold represents the size of price movements as a result of which a trader decides to place a new order. From this analysis, we conclude with the following two observations:

- In terms of the traders' profit objectives, the majority of traders place a new order based on small changes in the price. For that reason, based on a uniform distribution, the objective profit thresholds will be generated for each individual agent by drawing random numbers from a uniform narrow interval [0.01%, 2%]. The limits of such intervals are essential for reproducing the real FX market trading activity.

Algorithm 1 The core trading mechanism for the two groups of ZI-DCT0 agents

Require: initialise variables (event is upturn event, $x = p(t_0)$, Δx_{DC} (Fixed) ≥ 0)

```
if event is upturn event then
  if  $p(t) \leq x \times (1 - \Delta x_{DC})$ 
    event  $\leftarrow$  downturn event
     $x \leftarrow p(t)$ 
    Sell  $\rightarrow$  ZI-DCT0 contrarian agent, Buy  $\rightarrow$  ZI-DCT0 trend-following agent
  else
     $x \leftarrow \max(x, p(t))$ 
  end if
else //Event is downturn event
  if  $p(t) \geq x \times (1 + \Delta x_{DC})$  then
    event  $\leftarrow$  Upturn event
     $x \leftarrow p(t)$ 
    Buy  $\rightarrow$  ZI-DCT0 contrarian agent, Sell  $\rightarrow$  ZI-DCT0 trend-following agent
  else
     $x \leftarrow \min(x, p(t))$ 
  end if
endif
```

- The length of the traders' holdings on losing positions tends to be extensive. This indicates that many of the trading strategies of real traders involve a high degree of risk. This hypothesis was mapped onto the agent-based models of traders in which we defined the agents' risk appetites to be four times their profit objectives.

4.1.4 Trading strategy

The agents adopt the zero-intelligence directional-change event trading strategy (ZI-DCT0) (Aloud et al., 2011). The central idea behind ZI-DCT0 is that an agent will observe the changes in the price time series and commit to a fixed threshold based on intrinsic time rather than physical time. Physical time adopts a point-based system and is homogenous with time scales equally spaced based on the chosen time unit (second, minutes, etc.). In contrast, intrinsic time is irregularly-spaced in time as time triggers just at periodic events (Glattfelder et al., 2011). The intrinsic time basic unit is an *event* which is defined as the total price change exceeding a given fixed threshold set by the agent.

For studying an asset's price time series, given a fixed threshold of size Δx , the absolute price change between two local minimum and maximum prices is decomposed into a directional-change (DC) event of size Δx and its associated overshoot (OS) event (Glattfelder et al., 2011). The OS event represents the price movement beyond the fixed threshold Δx . A DC event is usually followed by a price OS event rather than a DC event in the reverse direction. A DC event takes one of the two forms - an upturn DC or a downturn DC event (Tsang, 2010). An upward run is a period between an upturn DC event and the next downturn DC event, while a downward run is a period between a downturn DC event and the next upturn DC event. An upturn (downturn) DC event ends a downward (upward) run, and starts an upward (downward) run.

Before the study of an asset's price time series, two variables are defined: the last high and low prices which are set to the asset's price at the start of the price sequence (Tsang, 2010). During an upward (downward) run, the last high (low) price is continuously updated to the maximum (minimum) of (a) the current price and (b) the last high (low) price. An upturn (downturn) DC event occurs once the absolute price change between the current price and the last low (high) price is higher (lower) than the threshold of size Δx .

Before the launch of the simulation, based on a continuous uniform distribution, each ZI-DCT0 agent is assigned a fixed threshold from a uniform interval [0.01%, 2%]. Two different equal groups of agents are identified based on ZI-DCT0: contrarian and trend-following agents. A contrarian agent opens a new position on the prediction that the price will move in the opposite direction. Trend-following agents take advantage of the price trend movement direction, assuming that the current trend will continue in the same direction. Algorithm 1 illustrates the core trading mechanism for these two groups of ZI-DCT0 agents.

Identifying a series of DC and OS events in the course of studying the price time series enables the agent to capture the short-term market dynamics involved based on their expectations of the market. Using the DC framework as the basis of the agent's strategy reduces the computational load and the complexity of price time series given that the DC framework yields a small number of price points for evaluation. Hence, the use of the DC framework allows an agent to detect fewer major periodic stable patterns enabling it to react to events and patterns of interest as specified by their strategy and threshold.

The threshold that an agent uses as part of the ZI-DCT0 strategy needs to be chosen with caution as it is essentially used to detect the market dynamics. Too high a threshold will mean that the agent is only able to detect very big changes in the prices (high or low) and may thus lose opportunities to gain from smaller movements in the market, whereas too low a threshold and very small variations in the price would be perceived as events and trigger the agent to consider acting.

A similar trading strategy in which the agents enter the market and operate in it only if the price fluctuations are above a certain threshold was introduced in (Alfi et al., 2009a). The authors suggest that introducing such a threshold, which modulates the number of active agents in the market, is one of the factors responsible for the exhibition of stylized facts in an ABM. The main difference between the strategy in (Alfi et al., 2009a) and the ZI-DCT0 is that the latter considers the direction and the overshoot of the price movement whereas the former does not.

The agents place buy and sell orders in the market and there are two types of orders: market orders and limit orders. Market orders are orders for immediate execution at the current price of a currency in the market. Hence, there is certainty that the market order will be executed, but uncertainty concerning the execution price. For a market order, an agent specifies the order size and type, with the order type being either a buy or a sell order. A position is opened once an agent has bought or short-sold any quantity of currency units, hence a position can be of type long (bought) or short (sold). The position volume stands for the number of bought or sold currency units executed in the market.

To acquire a comprehensive trading strategy for the agents, we incorporate two types of limit orders: profit-taking and stop-loss limit orders. For limit orders, a trader must declare the price at which it desires to close its current opened position. There is no guarantee that the declared price will be reached and as a result the execution of a limit order is not guaranteed. If a limit order is executed, there is a guarantee with regard to the execution price. The profit-taking limit order is an order to close a position once a certain profit threshold is reached to lock in a certain profit realisation. The profit-taking limit order is determined with reference to the agent's profit objective. For instance, for closing a long position the bidding price for the profit-taking limit order must be higher than the purchasing price, while for closing a short position, the asking price for the profit-taking limit order must be less than the selling price. The stop-loss limit order is an order to close a position once a loss threshold has been reached. Such a threshold is defined by the agent, subject to its risk appetite. For closing a long position, a stop-loss limit order becomes a market order after the bid price is at, or below, a specified loss threshold price (stop price) determined by the agent. On the other hand, for closing a short position, a stop-loss limit order becomes a market order after the ask price is at or above a specified stop price as defined by the agent. A stop-loss limit order seeks to limit the loss potential for a position. The closure of a position takes place as a result of either a margin call or an agent's decisions through market or limit orders explained in section 4.2.

An agent decides the type of an order based on the dynamics of the price time series. For opening a new position, an agent will decide on the market order. Hence, once the position is opened an agent will place two types of limit order: a profit-taking limit order to lock in a defined profit realisation and a stop-loss limit order to limit the loss potential for the position. For instance, a contrarian agent may open a new long position when the price falls by 0.04% and afterwards closes the position as soon as the price rises by 0.06%. By way of illustration, a contrarian agent opens a new short (long) position if the current bid (ask) price rises (declines) by a percentage which is equal to, or above, the threshold defined based on the agent's expectation of the market. A trend-following agent opens a new long position whilst the price is rising, whereas a short position is opened when the price is falling. Hence, a trend follower opens a new short (long) position if the current bid (ask) price declines (rises) by a percentage which is equal to, or above, the threshold used by the agent (refer to Algorithm 1).

4.1.5 Random trade

Market traders may place orders based on factors other than price movements (e.g. economic and political events). Hence, the price movements are influenced by exogenous variables. In our ABFXM, this is modelled by incorporating random trading into the agents' trading strategies. A random trade represents a trade placed beyond the notion of changes in price time series. When an agent's trading conditions are not met, it will place a random trade based on a probability equal to 10^{-6} which is set prior to the launch of the simulation.

4.1.6 Order size

The order size is an important element of our model since it is associated with the cautiousness of the agents. In the ABM literature, there are two main approaches to determining a single order size placed by an agent. The direct approach entails an agent's order size being equal to its wealth (Samanidou et al., 2007). This approach creates an extremely high trading risk, given that an agent invests its total cash in a single order. The second approach assumes that an agent's order size is a fraction of its cash (Martinez-Jaramillo and Tsang, 2009). The fraction value is determined by a constant trading proportion parameter.

In our ABFXM, the determination of the number of currency units of an order depends on the agent's available margin. If Units Available (UA) is the maximum number of currency units an agent a is allowed to trade based on its available margin, then the number of units that will be placed for an order is determined randomly within the range $[UA \times 10\%, UA]$. Experiments using the methods in (Martinez-Jaramillo and Tsang, 2009; Samanidou et al., 2007) have been carried out in Section 6.6.

4.1.7 Asynchronous trading time window

Timing is an important design issue in building ABMs. In an ABM, timing refers to many things including the time span of the previous information considered by the agents. The study by Daniel (2006) shows that the dynamics of the ABM can differ depending on whether trading is synchronous or asynchronous and concludes that the asynchronous nature of trading in financial markets essentially has to be considered in ABMs. Synchronous-time trading implies that at time t of the simulation run, all agents send their orders and update their beliefs. Afterwards, the time of the simulation run updates to $t+1$. When trading is asynchronous agents may place an order, hold or simply be inactive. As a result, the number of agents taking part in the total demand and supply is not constant at every single time of the simulation run.

We have modelled the asynchronous nature of trading and hence not all agents have to make a decision/trade at each time point since an agent does not depend on another agent's trading action to continue executing. We consider three important aspects: the agents' activation time in the market, the notion of market business hours, and holidays. As it is the case in real markets, not all agents are active at the launch of the simulation. We have set up an initial activation condition (INAC) to control when an agent will be active in the market for the first time. The INAC is based on a probability equal to 10^{-2} which is defined before the simulation starts.

Although the FX market operates on a 24-hour basis, the FX market centres' business hours play an important role on the level of transactions in the market. Hence, in our simulation, each agent is associated with local business trading hours corresponding to its geographical location. These are Sydney, Tokyo, London and New York. This means that although an agent can trade at any time during the day, the event of doing so outside its business hours is restricted based on a very small probability equal to 10^{-6} (except for executing limit orders).

4.2 Trading cycle and market clearance

A central feature of the simulation is how the trading cycle is sequenced, and how trading-related tasks are scheduled. In this section, an agent a 's trading cycle at time t is presented in terms of step structures:

Step 1. If an agent a has a limit order, the system will check its execution in terms of whether the limit order reaches the desired profit threshold or hits a loss threshold as explained in Section 4.1.4. The consequence is subject to two cases: (1) if the limit order is executable, go to step 7; (2) if the limit order is not executable, go to step 8.

Step 2. Based on a probability equal to 10^{-2} : (1) agent a will remain inactive in the market; (2) agent a will be active in the market, go to step 3.

Step 3. If time t is within agent a 's trading business hours, a performs step 5.

Step 4. If time t is outside agent a 's trading business hours, based on a probability equal to 10^{-6} , a either performs step 5 or will not participate in the market at time t .

Step 5. Agent a will take a decision to buy, sell or hold, based on its budget constraints and its trading strategy. The consequence is subject to two cases: (1) agent a places an order, go to step 7; (2) agent a takes a decision to hold since the price movement has not met a 's expectations, go to step 6.

Step 6. The price movements have not met the agent's expectations. Based on a probability equal to 10^{-6} , agent a decides whether to place a random order (go to step 7), or decides not to participate in the market at time t .

Step 7. The market-maker will execute the orders.

Step 8. Agent a 's portfolio will be updated.

The clearance mechanism of the ABFXM involves a very straightforward task. All the market orders at time t will be entirely executed, while limit orders will be executed when their constraints are satisfied. Subsequently, an update of the agent's portfolio will occur for each agent that has an executable order at time t . After that, the market's time turns from t to $t+1$. Accordingly, the bid and ask prices are adjusted to the prices at time $t+1$ using the currency pair dataset of the historical bid and ask prices. Based on the recent bid and ask prices, the portfolio will be updated for each agent that holds an open position at time $t+1$. Eventually each open position at time $t+1$ will be checked for a margin call, which is a procedure to close out the open position held by an agent as soon as the amount of cash in its account cascades under the minimum margin needed to cover the size of its currently open position. The intention of the margin call is to prevent an agent from losing more than the amount of cash available in the account.

5 Validation

In (Aloud et al., 2013), we established a set of stylized facts in the FX market-maker market by observing the behaviour of the traders and their transactions in a real HF dataset provided by OANDA. The identified stylized facts apply on transactions data and can be grouped under three headings: seasonality, correlation behaviour and scaling laws. These stylized facts describe the transactions data in the FX market from different angles. In addition to confirming well-known stylised facts as observed in other financial markets, we have established a new set of stylised facts for the FX market. We use these stylized facts to validate and confirm that the ensuing behaviour from the modelled ABFXM in terms of transactions data is approximating that of the real market. In essence, we use the stylized facts as the means to validate that the ABMFX is a model of the real market and is able to approximate the real market behaviour to a satisfactory extent.

5.1 Dataset

Our study is supplied with two high-frequency datasets by OANDA. OANDA offers an online foreign currencies trading platform and is considered to be the largest such online platform in terms of trade volumes and liquidity. The first dataset comprises unique high-frequency data 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 contains about 147 million historical transactions carried out by 45,845 different accounts trading in 48 different currency pairs under the same terms and conditions. Each transaction contains the following information: the transaction type, the transaction timestamp, the traded currency pair, the execution price, the units and the amount traded. We have used this dataset to establish stylized facts of the transactions in FX markets (Aloud et al., 2013). The second dataset comprises 2.25 years data samples of EUR/USD tick-by-tick market data from January, 1 2007 to March, 5 2009. The data take the form of (a) a bid and (b) an ask price of a currency pair at (c) a timestamp.

The simulation uses historical bid and ask prices for the EUR/USD currency pair over a defined one month period by feeding these prices into the market and having the agents respond to the price changes. The data are fed through the market-maker component of the simulation. We use this dataset as the purpose of this work is to develop an ABFXM that is able to approximate the stylized facts of transactions that are exhibited in the real FX market data and explore their origins. These stylized facts relate to transactions data (a transaction represents an executable order in the market) and hence the price data are used to anchor the simulation results so that we can establish whether the transactions data generated by the ABFXM have the same characteristics and properties

Parameter	Value
Number of trading periods	3.00E+06
Number of agents	10^4
Set of leverage	{1, 5, 10, 20, 25, 30, 40, 50}
Set of FX trading session time zones	{Sydney, Tokyo, London, New York}
Threshold range of agents' profit objectives	[0.01%, 2%]
Initialize cash range for a power law distribution	$[10^3, 10^9]$
Scaling exponent for a power law distribution	1.2

Table 1: Initial market parameters configuration in the simulation.




Figure 1: Intraday seasonality of EUR/USD trade numbers and volumes from (a) the real data and (b) the simulation results.

as the real market data. The historical dataset of prices acts as a constant so that multiple simulation runs can be performed in which aspects of the trading behaviour of the agents can be modified and systematically studied without having to worry about price setting in the market. This enables us to study the ensuing behaviour and data generated from the market against those of the real market and also compare different settings against one another.

5.2 Results

The results are averaged over 30 independent simulations, each conducted over a different time period of one month with different initial seeds of the random number generators. The one month time period is sampled randomly from the EUR/USD historical price dataset which comprises 2.25 years data samples of EUR/USD tick-by-tick market data. We run 30 simulations with the same parameters configuration values, but with different seeds and a price dataset of one month, to ensure that the results of the simulation are constant; this allows an assessment of the robustness and accuracy of the simulation results. The market and agent parameters are as explained in the previous sections and are summarised in Table 1¹.

Seasonality. The FX market data exhibit periodic patterns called seasonality which have been confirmed in studies including (Dacorogna et al., 2001; Ito and Hashimoto, 2006; Aloud et al., 2013). Using a uniform time grid with 24 hourly intervals, the intraday seasonality relates quantities of the trading activities (trade numbers and volumes) to the time of day when these are observed. The shape of the FX market intraday seasonality is referred to as a double U-shape or camel shape because of the two peaks in which the second peak is higher than the first. This indicates that the transactions start to peak within the opening business hours of the main FX market centres in the morning, decline during the lunch break and then peak in the afternoon. They subsequently gradually decline outside the main market centres' business hours. Figure 1 shows the intraday seasonality of transactions from (a) the real data and (b) the ABFXM. The ABFXM approximates the seasonality of the transactions in the FX market and the generated transactions clearly preserve the double U-shape (camel-shape) pattern in which the second peak is higher than the first.

A paired t-test is used to compare the means of two sets of seasonal data: (a) the real FX market data and (b) the simulation data. The seasonality observations from the FX market data are paired with the seasonality observations from the ABFXM. The size of the two sets is equal, denoted by n , where

¹A sensitivity analysis of the impact of these parameters in reproducing realistic agent-based models of FX market traders' behaviours is conducted in Section 6.

	<i>tsv</i>	corr
Trade numbers	2.03E-10	+ 0.90
Trade volumes	1.11E-10	+ 0.86
Number of opening positions	2.02E-10	+ 0.91
Number of closing positions	-6.08E-10	+ 0.91

Table 2: T-statistic values and standard errors - Intraday seasonality statistics.

	(a)	(b)
Trade numbers and volumes	+ 0.81	+ 0.92
Numbers of buy and sell orders	+ 0.95	+ 0.99
Numbers of opening and closing positions	+ 0.98	+ 0.98
Intraday price volatilities and trader numbers	+ 0.43	+ 0.46
Intraday price volatilities and volumes	+ 0.28	+ 0.31

Table 3: Correlation coefficients computed for the different EUR/USD transactions from (a) the real data and (b) the simulation results.

n=24 for the intraday seasonality while n=168 for the intraweek seasonality. For the paired t-test to be valid, the differences between the seasonality of the FX market data and the ABFXM data need to be approximately normally distributed. Accordingly, a normality test has to be conducted. In terms of intraday seasonality, we can test the normality using a normal distribution plot, since the data size is less than 30. In contrast, for the intraweek seasonality we test for normality using a normal distribution plot and a Jarque–Bera test.

According to the paired t-test related to the hypothesis ($\mu_R - \mu_S = 0$), **the simulation reproduces the intraday seasonality observed in real FX market with a 95% confidence level.** Table 2 shows the intraday seasonality statistics, the t-statistic values and Pearson correlation values, where the t-critical value two-tail = 2.07.

Correlation. The correlation stylized facts concern the positive linear relationship that was found between different quantities of trading activities (Aloud et al., 2013). The simulation results exhibit very similar correlation behaviour to that of the observed FX market transactions as reported in Table 3. There is a positive linear relationship between (i) trade numbers and volumes, (ii) numbers of buy and sell executed orders, (iii) numbers of opening and closing positions, and (iv) intraday price volatilities and trade numbers and volumes.

Scaling laws. Scaling laws establish quantitative relationships between the size of price movements and the market trading activities as a function of the time interval at which market transactions are measured. To define scaling laws for the EUR/USD transactions, and using the directional-change event approach (Section 4.1.4), we identify directional-change (DC) and overshoot (OS) events in the EUR/USD price time series for thresholds of different sizes, ranging from 0.10% to 0.80%. Given a fixed percentage threshold, the average observed (a) trade numbers, (b) trade volumes, number of (c) opening and (d) closing positions observed during a DC and an OS event, is scale-invariant to the size of this threshold (Aloud et al., 2013). Our study has also uncovered six quantitative relationships amongst the four scaling laws, holding across EUR/USD transactions. These relationships specify that, on average, an OS event contains roughly twice as many trade numbers and volumes, and numbers of opening and closing positions, as a DC event. In addition, an OS event contains approximately the same number of opening and closing positions, while similar features hold for a DC event. We examine the exhibition of the scaling laws by defining the linear dependence. The adjusted R^2 value estimates the significance of fit of the line and can take any value less than or equal to 1, with a value closer to 1 signifying a better fit. Consequently, the closer the points to the line, the better the fit. Note that the adjusted R^2 value is not sensitive to the number of observed points within the data sample.

The four scaling laws are plotted in Figure 2 for (a) the real data and (b) the simulation results. Table 4 shows the adjusted R^2 values and standard errors of the fits. The standard errors of the fits are roughly equivalent to the estimated errors for scaling laws in the real data. For the adjusted R^2 values, it is worth noting that these values are close to 1 which implies a good fit. The simulation results also exhibit the six quantitative relationships amongst the four scaling laws (Figure 2 (b)).




Figure 2: Scaling laws are plotted for (a) the real data and (b) the simulation results.

	(a)				(b)			
	DC		OS		DC		OS	
	R^2	SE	R^2	SE	R^2	SE	R^2	SE
Trade numbers	0.90	0.37	0.90	0.35	0.91	0.37	0.90	0.31
Trade volumes	0.91	0.38	0.90	0.35	0.94	0.35	0.94	0.31
Opening positions	0.90	0.35	0.91	0.35	0.92	0.35	0.93	0.31
Closing positions	0.90	0.36	0.90	0.35	0.91	0.39	0.90	0.35

Table 4: The adjusted R^2 values of the fits, plus their standard errors (SE) for the scaling laws relationships measured from the simulation outcome under directional-change (DC) and overshoot (OS) events.

Case	Value	Seasonality		Correlation	Scaling laws
		Transactions	Volumes		
Number of agents	10 ²	✓	☒	☒	☒
	10 ³	✓	✓	☒	☒
	10 ⁴	✓	✓	✓	✓
	10 ⁵	✓	☒	☒	☒
Initial activation condition	On	✓	✓	✓	✓
	Off	☒	☒	☒	☒
Limit orders	On	✓	✓	✓	✓
	Off	☒	☒	☒	☒
Initial wealth distribution	Power law distribution	✓	✓	✓	✓
	Same wealth	✓	✓	☒	☒
	Continuous uniform	✓	☒	☒	☒
Profit objective	Normal distribution	☒	☒	☒	☒
	Continuous uniform	✓	✓	✓	✓
	Power-law distribution	☒	☒	☒	☒
Order size	Equal to wealth	✓	☒	☒	☒
	Fraction of wealth	✓	✓	☒	☒
	Available margin	✓	✓	✓	✓
Trading time window	On	✓	✓	✓	✓
	Off	☒	☒	☒	☒
Risk Appetites	Extremely high	✓	✓	☒	☒
	High	✓	✓	✓	✓
	Low	☒	✓	☒	☒
	Zero	☒	✓	☒	☒
	Mixed risk levels (2-4)	☒	✓	☒	☒
Trading Strategy	50%-50% contrarian and trend-following	✓	✓	✓	✓
	Just contrarian	✓	✓	☒	☒
	Just trend-following	✓	✓	☒	☒
	40% contrarian and 60% trend-following	✓	✓	☒	✓
	70% contrarian and 30% trend-following	☒	☒	✓	✓
	Random trades disabled	✓	✓	☒	☒

Table 5: Exhibition of the FX market stylized facts of transactions data for each group of experiments.

6 Exploring the origin of the FX market stylized facts

Having established a model of the FX market that is able to satisfactorily approximate the observed stylized facts, we subsequently perform an exploration of the ABFXM's constituent elements in order to study their impact on the emergence of the stylized facts. For each relevant element, we performed 10 runs each of which lasted one simulated month. We report the average of the obtained results and draw conclusions regarding their impact on the market transactions. Table 5 summarizes the results of the emergence of the stylized facts on transactions data for each group of experiments. **The identified stylized facts of the FX transactions are grouped under three headings: seasonality, correlation behaviour and scaling laws. These stylized facts describe the transactions data in the FX market. Therefore, all these stylized facts have to emerge in an ABM in order to confirm that an ABM is able to approximate the real market behaviour to a satisfactory extent. Hence, these stylized facts are used to validate and confirm that the resulting behaviour from the ABFXM in terms of transactions data is approximating that of the real market. For more details refer to the appendix.**

6.1 Number of agents

The FX market involves a large number of participants and we suppose that the total number of agents N in an ABM may be critical in reproducing the stylized facts. To examine the impact of the number of agents on the dynamics of the market behaviour, we performed several experiments in which we scaled N from 10^2 to 10^5 .

As per the results, the full set of stylized facts is obtained only for a finite value of $N = 10^4$. This is due to a small N in the market affecting the intensity of transactions. A large number of agents affects the equilibrium of demand and supply in the market in terms of number and volume of transactions, which may lead to the flow of transactions moving in one direction, and essentially affecting the appearance of the market stylized facts. Other studies have shown that the number of agents in the market can have an impact. For instance, in (Egenter et al., 1999), the impact of the total number of agents has been shown for the Kim and Markowitz and the Lux and Marchesi models. In these two models, the stylized facts appear for a certain range of the total number of agents, and in most cases the stylized facts do not appear for a large number of agents. In the Lux and Marchesi model, the total number of agents affects the excess demand, and therefore affects the appearance of the stylized facts of the price time series. In (Yeh, 2001), the study shows that the stylized facts tend to vanish when the total number of agents increases in the Artificial Stock Market under consideration.

6.2 Initial activation condition

In our ABFXM and at the start of the simulation, not all agents are active in the market. We have modelled this as an initial activation condition (INAC) which is, in essence, a probability equal to 10^{-2} . This is to emulate the fact that in real life, not all agents are active. To examine the impact of INAC on the stylized facts, we switch it off and hence all agents are active right from the start. Our results suggest that the incorporation of the INAC is actually quite important for the emergence of stylized facts in the ABFXM. To the best of our knowledge none of the existing ABMs in the literature imply a constraint with regard to the agents' initial activation time in the market. ABMs typically assume that all agents are active in the market from the start of the simulation.

6.3 Limit orders

The generation of the two types of limit orders in the market – profit-taking and stop-loss limit orders – has an effect on the emergence of stylized facts. The reported results in this section confirm that the generation of limit orders represents an important element in reproducing the transactions in the FX market. The inclusion of limit orders endows agents with a comprehensive trading strategy through which they can lock a certain profit realization, or seek to limit the potential losses with regard to their open positions. The absence of limit orders suggests that the agent strategy is stripped away from an essential element. These results confirm the importance of the generation of limit orders in an ABM in line with the studies in (Daniel, 2006; Martinez-Jaramillo and Tsang, 2009; Martinez-Jaramillo, 2007).

6.4 Initial wealth distribution

The agents in the ABFXM are endowed with different amounts of wealth based on a power law distribution. To assess the impact of different wealth distribution methods we performed experiments using two different settings: (1) all agents are endowed with the same wealth as in (Martinez-Jaramillo and Tsang, 2009) and (2) agents are endowed with different amounts of wealth, using a continuous uniform distribution. The results show that the correlation between price volatility and market transactions is exhibited when agents at the launch of the simulation possess different wealth by means of a power law distribution.

According to the results, we would argue that the assumption that agents possess the same wealth at the launch of the market as has been done by (Martinez-Jaramillo and Tsang, 2009) may not be representative of real market conditions. This is because real markets involve heterogeneity in the wealth distribution of the traders. In the literature, a number of studies found that the income distribution in numerous countries follows a power-law distribution (Pierre, 1965; Persky, 1992). Following on from these studies, we quantify the income distribution in the markets by means of the income distribution of a country's population. Asset price and trade volumes evolve endogenously according to the imbalance of demand and supply in the market. This suggests an unfairly distributed initial wealth for the agents based on a power-law distribution.

6.5 Distribution of the profit objective

The precise setting for the agent's profit objective in an ABM is one of the important elements involved in reproducing the stylized facts of the HF FX market. The choice of continuous uniform distribution for the agents' profit objectives, is capable of approximating the stylized facts. To investigate the impact of different distribution methods of the agents' profit objectives on the stylized facts, we conducted experiments in which we used a power-law and a normal distribution as alternatives.

The results emphasise that the price dynamics in the market are a result of the existence of heterogeneous expectations amongst traders. The power-law and the normal distributions do not appear to characterise adequately the profit objective values. Consequently, this results in a large group of agents having similar behaviour and preferences whereas, in reality, traders have a variety of preferences and demands.

6.6 Order size

In this section, we examine the impact of adopting two approaches that have been used in the literature to determine an agent's order size on the ensuing market transactions and stylized facts. According to the first one, the order size is equal to the agent's wealth (Samanidou et al., 2007). The second approach uses a constant trading proportion parameter that controls the size of an agent's order in terms of the agent's current wealth (LeBaron, 2001; Raberto and Cincotti, 2005; Martinez-Jaramillo and Tsang, 2009). With regard to this second approach, we performed several experiments in which we scaled the trading proportion parameter from 10% to 90%. Due to space limitations and the insignificant changes in reporting all the results, we only report those associated with having the trading proportion parameter equal to 25%, 50% and 75%. For values under 10% the changes were insignificant, while for values above 90%, the results resemble those obtained from using the first approach in which the amount that an agent invests in the market is proportional to the agent's wealth.

Assuming that a single order size placed by an agent in the market is equal to the agent's wealth, such as in (Samanidou et al., 2007), does not reproduce the stylized facts. This is because it creates an extremely high trading risk, since an agent invests the full cash it acquires in a single order. This in turn triggers cascaded margin calls, which induce a price fall. A second reason is the lack of heterogeneity in the amount of an agent's investments during the simulation. The second approach entails that the amount the agent invests in the market will be a fraction of its cash, based on a constant trading proportion parameter (LeBaron, 2001; Martinez-Jaramillo and Tsang, 2009). Similarly to the first one, this second approach induces a lack of heterogeneity of investments in the market. In addition, having a high value in terms of the trading proportion parameter causes cascading margin calls in the market. In contrast, having small values in terms of the trading proportion parameter results in drastic changes in the stylized facts. This is due to the low liquidity in the market and the lack of correlation between the trade numbers and volumes.

6.7 Trading time window

Although the FX market operates 24 hours a day and does not have fixed trading time windows as the stock market, the FX market centres' business hours around the world have a significant effect on characterizing the daily flow of the market's transactions. The intensity of transactions increases during the opening hours of the FX market centres in the morning, whereas it decreases during the closing hours. The overlapping business hours of two or more FX market centres witnesses a high level of transactions compared to the other times.

Initially, we decided to model the agents without incorporating the role of FX market trading sessions and holidays. The motivation behind our decision was to account for the significant impact of algorithmic trading. Algorithmic trading, also referred to as automated trading, uses computer programs to place orders automatically which may be outside a trader's normal business hours. It has attracted many FX traders and, most importantly, is one of the reasons for the growth of the FX market. However, not modelling business hours has a significant impact on the emergence of stylized facts.

The results present evidence of a strong correlation between market transactions and the time of day. The correlation between market transactions and volatility requires us to model a heterogeneous trading time window with the help of information related to the main FX market centres' business hours around the world.

6.8 Risk Appetites

The risk appetite parameter is used to model the cautiousness and aggressiveness of agents with regard to trading. From the microscopic analysis of the OANDA individual traders' historical transactions, we found on average, that the traders' risk appetites are higher than their profit objectives (as explained in Section 4.1.3). Accordingly, in our ABFXM, we defined the agents' risk appetites to be four times that of their profit objectives.

To test the impact of such an element on reproducing the stylized facts, we varied the risk appetites' parameter in the simulation. Although we have conducted a number of experiments, we report the results of five indicative cases in which the agents' risk appetites relate to their profit objectives as follows:

1. Extremely high risk: the agents' risk appetite is higher than four times their profit objectives.
2. High-risk: the agent's risk appetite is defined as being four times that of its profit objective. The high risk appetite setting is the default value used in our ABFXM.
3. Low-risk: the agent's risk appetite is equal to the agent's profit objective.
4. Zero-risk: an agent will close its open position once the unrealized losses for the position fall below zero, which means that the position is in loss.
5. Mixed: prior to the launch of the simulation, the agents are assigned randomly a risk appetite among the three risk appetites 2-4 as described above.

The level of the agent's risk appetite is considered to be an essential element in an ABM. In our agent-based models of traders, an agent's risk appetite is articulated in terms of the agent's profit objective. We found that setting a high level of agents' risk appetites compared to the agents' profit objectives means that the ABFXM is capable of reproducing, to a certain extent, the trading activity in the market. However, having the agents' risk appetite higher than four times their profit objectives, results in negative correlations between the intraday price volatilities and trade numbers (- 0.02) and volumes (- 0.03). This is because of the implications of high risk activity on an agent's wealth performance, which triggers margin calls. The right setting of the agents' risk appetites parameter significantly affects the correlation between intraday price volatility and market trading activities. In particular, using a low risk appetite results in a negative correlation between intraday price volatility and market trading activities. Such a negative correlation may be affected by the high intensity of trading activities when the price collapses. The other risk settings also did not seem to be able to reproduce the stylized facts of the FX market. Accordingly, the correlation results confirm what has been observed in market data: traders in the high-frequency market tend to hold losing positions for a long period of time.

6.9 Contrarian and Trend-Following Strategies

In our ABFXM, we consider a population whose trading strategies are organized into two major classes with equal number of agents: trend-following and contrarian ZI-DCT0 trading strategies. We aim to study the impact of using just one of these strategies for the entire population and hence we conducted experiments where the adopted trading strategy in the market was either the trend-following or the contrarian strategy.

Also, we examined the effect of the existence of different percentages of contrarian agents in comparison to trend-following agents on the exhibition of the stylized facts. Based on the results obtained, the examination can be divided into two cases, with each case involving a series of experiments. In the first case, the number of contrarian agents is between 30% - 70% of the total population. In the second case, the number of contrarian agents is more than (i) 70% or (ii) less than 30% of the total population: these two sub-cases generate very similar results due to a very large number of agents having the same trading strategy. Hence, we only present an example of the latter in this section due to the similarity of the results and for brevity.

In addition, a series of experiments were conducted to study the impact of the inclusion of random trades in the market; random trades have been incorporated in the ABFXM to emulate the fact that in reality traders take decisions influenced by exogenous to the market events.

In our ABFXM, we limit the trading strategies to two classes of trend-following and contrarian trading strategies, in which agents can enter or leave the market, depending on the price behaviour. This is an extreme simplification with regards to modelling the market traders' behaviour. The results from our ABFXM show that the simplification in the two classes of trading strategies is enough to achieve an understanding, to a certain

extent, of the original market trading activity stylized facts. However, the possibility of a wider class of trading strategies in an ABFXM, and the traders' decisions on the basis of previous performance, is without doubt an interesting area with regard to future exploration. The problem, in this respect, is that the analysis and exploration of trading activities in the ABFXM can become highly misleading and difficult to study, because the broader class of trading strategies may produce a source of complexity for the market analysis.

Therefore, the results support that competition between trend-following agents and contrarian agents would have a tendency to stabilize the market. The key role played by the coexistence of such agents' trading strategies is one of the main elements responsible for the exhibition of the stylized facts in an ABFXM. Random trading governs the likelihood of the agents changing their trading strategy, and it is a vital element of the model. In any case, the intraday seasonality statistics of the different market trading activities are exhibited in the limit where the agents' population is composed only of trend-following or contrarian traders, even without random trades taking place. Therefore, these simple experiments show us which stylized facts depend critically on the interaction between different trading strategies. The generalized scaling law properties and the correlation behaviour results appear to be sensitive to the heterogeneity of agent trading strategies.

Similar studies by (Alfi et al., 2009a,b,c) showed that the stylized facts are exhibited in an ABM by introducing the possibility of the agents entering or leaving the market depending on the behaviour of the price. The mechanism to enter or leave the market is linked to the existence of a threshold that the agents use to decide to operate in the market by comparing the threshold with the price movement.

7 Conclusions

Our work has been motivated by the need to have a comprehensive understanding of electronic market workings and the forces that drive the flow of the market trading activity. This is an essential prerequisite for the development of computational-intelligence based automated strategies for electronic markets. Our particular focus is the FX market which is the market for the trading of currencies.

In this paper, we described an agent-based model of the FX market which is able to approximate the stylized facts of the real FX market transactions data (Aloud et al., 2013). The contribution of our work is fivefold. Firstly, we have developed an agent-based model of the FX market (ABFXM). To the best of our knowledge, there are no other works in the literature that have modelled the FX market using an agent-based approach. An agent-based modelling approach is particularly suited for modelling complex systems with multiple entities, each governed by its own characteristics and behaviour and whose interactions give rise to the overall system behaviour and phenomena that are not under the direct control of just one of the agents. Agent-based models enable us to study complex system behaviour in a systematic way, something that analytical and more traditional approaches to studying economic systems behaviour are not suited for.

Our second contribution is in the construction of the agent-based model of the FX market itself and in particular, the modelling of heterogeneity and asynchronous trading. One of the fundamental and driving hypothesis of our work is that traders are non-homogeneous. Heterogeneity is inherent in real markets, so in creating ABMs this is an important characteristic that needs to be modelled. Hence we have tried to model heterogeneity through a number of ways. Heterogeneity in our ABM is modelled in the agents' initial wealth, their trading strategies, their profit objectives, their risk appetites and their trading windows which represent geographical distribution. Furthermore, we have constructed a new approach for characterising an agent's order size which is dependent on the agent's available margin and determined randomly within a specified range. This approach provides a further source of heterogeneity in the agents' trading activities and hence the market. This is unlike other works in the literature such as (i) the work reported in (Samanidou et al., 2007) where the agent's order size is equivalent to the agent's total wealth; and (ii) the approach adopted in (LeBaron, 2001; Raberto and Cincotti, 2005; Martinez-Jaramillo and Tsang, 2009) where an agent's order size is a fraction of its cash based on a constant trading proportion parameter. These two approaches for modelling the agent's order size entail lack of heterogeneity in the amount of an agent's investments during the market run. Our results suggest that neither of them are able to reproduce the full set of stylized facts. Furthermore, the first approach generates an extremely high trading risk, as an agent invests the full cash it acquires in a single order. A high parameter value in the second approach also generates a high trading risk. This sequentially triggers cascaded margin calls, which induce a price fall in the market and affect the agent's performance. Whereas there are other works that have modelled heterogeneity such as (Martinez-Jaramillo and Tsang, 2009; Yang, 2002; Arthur et al., 1997),

these do so in a limited way through a small number of elements. No other work in the literature has used that many and as varying elements of heterogeneity as we have.

Another feature which is novel in our ABFXM, is the modelling of the asynchronous nature of trading as exhibited in real markets. To this end, we have considered four important factors: the agents' activation time in the market, the notion of market business hours, holidays and the introduction of a threshold to trigger the agents' activities. Again this is in contrast with what has been done in the literature so far in which some ABMs imply synchronous-time trading which means that at time t of the market each agent updates their beliefs and accordingly places an order (Jefferies et al., 2000; Martinez-Jaramillo and Tsang, 2009). Synchronous-time trading in an ABM induces a high computational cost as at every time point each agent has to carry out some computation. The computational cost should not be ignored in the design and the construction of an agent-based financial market. The Genoa Artificial Stock Market (Raberto and Cincotti, 2005) determines randomly that an individual agent will be active in the next period of the market. However, such randomisation is not adequate, as it could result in ambiguity regarding the arrival time of an order.

The third contribution of our work comprises the use of a strategy which is based on the directional change event approach (Aloud et al., 2011) and is able to identify and respond to periodic patterns in the price time series. The agents have been modelled as simple ZI-DCT0 agents that commit themselves to a fixed threshold (but this is different for each agent) to detect periodic fixed patterns in the price time series. This is different from other works in the literature (Alfi et al., 2009a; Martinez-Jaramillo and Tsang, 2009; LeBaron et al., 1999; Gode and Sunder, 1993). The ZI-DCT0 agents avoid the complexity of intelligent strategic behaviour and also the vagueness of zero-intelligence behaviour in which agents trade randomly subject to budget constraints. We show that the stylized facts of real markets can be reproduced by ZI-DCT0 agents, subject to budget constraints.

The fourth contribution is in the use of the ABFXM to explore and study the origins of the stylized facts in the FX market. We have undertaken a comprehensive exploration of the constituent elements of the ABFXM and their impact on the stylized facts of transactions data and have identified a set of sufficient conditions which result in the emergence of the stylized facts similarly to the real market data. Our study suggests that among the key elements that affect the emergence of the transactions stylized facts are the number of agents in the market and heterogeneity in the agents' strategies which in our model takes a number of forms: varying profit objectives, risk appetites and wealth distribution. The generation of limit orders also seems to play an important role as this provides another dimension to the agent's trading strategy coupled with orders of varying size. The mixture of contrarian and trend-following ZI-DCT0 agents in the market also seem to have an impact on the emergence of the stylized facts.

The fifth contribution is in the use of a unique high frequency dataset of historical transactions data to undertake this study. We have been able to study in a systematic way the effect of each of the constituent elements of the market on the stylized facts through anchoring the simulation results by using this historical prices dataset. The use of this dataset enables us to run multiple simulation runs and draw comparisons and conclusions for each market setting against the real data but also compare different market settings against each other. As far as we know, this has not been done before and in essence serves as an essential validation step in our understanding of what drives market behaviour and the emergence of the stylized facts in the FX market.

Our work is different from other works in financial markets such as (Alfi et al., 2009a; Martinez-Jaramillo and Tsang, 2009; Daniel, 2006). We study in depth the constructed ABFXM and focus on performing a systematic exploration of the constituent elements of the ABFXM and their impact on the dynamics of the FX market behaviour. In particular, our study explores and identifies a set of sufficient conditions under which the stylized facts of the FX market trading activity are exhibited in the ABFXM. The key elements are the ZI-DCT0 agents, heterogeneity which has been embedded in the ABFXM in different ways, asynchronous trading time windows, initial activation conditions and the generation of limit orders. We also show that the dynamics of the market trading activity depend on the number of agents one considers.

One of the criticisms of ABMs is the complexity in their setup and calibration. In our work, we have attempted to offer as simple a description of the elements involved in the ABFXM that are capable of reproducing, to a certain extent, the stylized facts of the HF FX market transactions. This has important consequences for investigating the origins of FX market transactions' stylized facts by means of exploring each of the model's elements in isolation and their impact on the other elements. The simplicity of our ABFXM relies on two main aspects: the use of a simple market mechanism in which we used currency pair historical prices as a substitute for generating artificial prices from the ABM, and the simple repre-

sensation of the trading agents. The agents' trading strategy is based on the ZI-DCT0 trading strategy in which the agents respond by buying or selling to a periodic pattern in the price time series based on their expectations of the market (Section 4.1.4). The use of the unique HF dataset of FX market transactions enables us on the one hand to calibrate our model and on the other to validate our approach.

There are several avenues for future work and one of them is studying the price formation on the FX market which is organized as market-makers markets. As the purpose of this work was to understand the workings of the FX market and develop an agent-based model of the market able to approximate the FX stylized facts, the market-maker is currently using the historical dataset to quote prices. One line of future research is to remove this feed and allow the market-maker to generate artificial prices and explore and study the characteristics of the price time series generated. We have been experimenting with different agent strategies such as for instance the Zero Intelligence (ZI) and Genetic Programming (GP) based strategies as part of a study to assess the level of intelligence required to reproduce the stylized facts in the FX market (Aloud and Fasli, 2016). These results indicate that complex strategies may not be necessary for the reproduction of the stylized facts, but simpler strategies such as the ZI-DCT0 used for this study may be more appropriate. To this end, we wish to explore further computational intelligence techniques and their impact on the emergence of the stylized facts. Another avenue for future work is to enable dynamic thresholds within the ZI-DCT0 strategy for detecting arbitrage opportunities in a price time series. The ABFXM could also be extended in several directions, such as for instance by introducing further assets to explore more complex trading behaviours.

The constructed ABFXM can be used as an experimental laboratory for exploring and identifying conditions as a result of which collapses occur in the market. Simulations have the potential to answer essential questions such as: Under what conditions do price collapses occur? Can we identify early signs of a market crisis? How can the market rules be improved to avoid excessive price collapses? A mechanism for identifying the potential signs of financial crises could be developed, along with a means by which we can analyse patterns with regard to what has happened prior to a crisis. An early warning system for the financial market could offer the means to identifying early indications of a financial crisis and help prevent enormous losses, improving the overall market efficiency.

Acknowledgments

This research was funded by the Research Center of Humanities, Deanship of Scientific Research, King Saud University. The authors would like to thank OANDA Corporation for providing the FX market high-frequency datasets.

References

- A. Booth, E. Gerding, F. McGroarty, Automated trading with performance weighted random forests and seasonality, *Expert Systems with Applications* 41 (2014) 3651–3661.
- M. Fasli, *Agent technology for E-commerce*, Wiley and Sons, 2007.
- M. Fasli, M. Michalakopoulos, e-Game: A platform for developing auction-based market simulations, *Decision Support Systems* 44 (2) (2008) 469–481.
- M. Dacorogna, R. Gençay, U. Müller, R. Olsen, O. Pictet, *An introduction to high-frequency finance*, Academic Press, San Diego, 2001.
- S. Martinez-Jaramillo, E. Tsang, An heterogeneous, endogenous and co-evolutionary GP-based financial market, *IEEE Transactions on Evolutionary Computation* 13 (1) (2009) 33–55.
- M. Aloud, M. Fasli, E. Tsang, A. D. R. Olsen, Stylized facts of trading activity in the high frequency FX market: An Empirical Study, *Journal of Finance and Investment Analysis* 2 (4) (2013) 145–183.
- R. Verma, K. B. Kenneth, *Operations and Supply Chain Management: World Class Theory and Practice*, Cengage Learning, 2010.

- M. Evans, R. Lyons, Exchange rate fundamentals and order flow, Tech. Rep. 13151, NBER, 2004.
- 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.
- B. LeBaron, A builder's guide to agent based financial markets, *Quantitative Finance* 1 (2) (2001) 254–261.
- W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, P. Tayler, Asset pricing under endogenous expectations in an artificial stock market, in: W. Arthur, D. Lane, S. Durlauf (Eds.), *The economy as an evolving, complex system II*, Addison Wesley, Redwood City, CA, 15–44, 1997.
- S.-H. Chen, C.-H. Yeh, Evolving traders and the business school with genetic programming: a new architecture of the agent-based artificial stock market, *Journal of Economic Dynamics and Control*, Elsevier 25 (3-4) (2001) 363–393.
- M. Levy, S. Solomon, Dynamical explanation for the emergence of power law in a stock market model, *International Journal of Modern Physics C* 7 (1996) 65–72.
- 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.
- S. Sunder, Market as an artifact aggregate efficiency from zero intelligence traders, in: M. Augier, J. March (Eds.), *Models of a Man: Essays in memory of Herbert A. Simon*, MIT Press, Cambridge, 501–519, 2004.
- TAC, Trading Agent Competition, See <http://tradingagents.eecs.umich.edu>, 2016.
- G. Daniel, Asynchronous simulations of a limit order book, Ph.D. thesis, University of Manchester, 2006.
- V. Alfi, M. Cristelli, L. Pietronero, A. Zaccaria, Minimal agent based model for financial markets I: origin and self-organization of stylized fatcs, *The European Physical Journal B* 67 (3) (2009a) 385–397.
- B. LeBaron, Agent-based Computational Finance, in: L. Tesfatsion, K. L. Judd (Eds.), *Handbook of Computational Economics*, vol. 2, chap. 24, Elsevier, 1 edn., 1187–1233, 2006.
- C. Hommes, Heterogeneous Agent Models in Economics and Finance, in: L. Tesfatsion, K. L. Judd (Eds.), *Handbook of Computational Economics*, vol. 2 of *Handbook of Computational Economics*, chap. 23, Elsevier, 1109–1186, 2006.
- E. Samanidou, E. Zschischang, D. Stauffer, T. Lux, Agent-based models of financial markets, *Reports on Progress in Physics* 70 (2007) 409–450.
- M. Aloud, E. Tsang, R. Olsen, A. Dupuis, A directional-change events approach for studying financial time series, *Economics Papers* No 2011-28.
- M. Levy, H. Levy, S. Solomon, A microscopic model of the stock market: cycles, booms and crashes, *Economics Letters* 45 (1) (1994) 103–111.
- B. Pierre, Pareto (Vilfredo) - Cours d'économie politique, *Revue Économique*, Programme National Persée 16 (5) (1965) 811–812.
- J. Steindl, *Random processes and the growth of firms - a study of the pareto law*, Charles Griffin and Company, London, 1965.
- B. LeBaron, Building the Santa Fe Artificial Stock Market, Tech. Rep., Brandeis University, 2002.
- G. Kim, H. Markowitz, Investment rules, margin, and market volatility, *Journal of Portfolio Management* 16 (1) (1989) 45–52.
- 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.

- E. Tsang, Directional changes, definitions, Tech. Rep. 050-10, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, UK, 2010.
- 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.
- E. Egener, T. Lux, D. Stauffer, Finite-size effects in Monte Carlo simulations of two stock market models, *Physica A* 268 (1999) 250–256.
- C. H. Yeh, The influence of market size in an artificial stock market: the approach based on genetic programming, I-Shou University, Kaohsiung, Taiwan, 2001.
- S. Martinez-Jaramillo, Artificial financial markets: an agent based approach to reproduce stylized facts and to study the Red Queen Effect, Ph.D. thesis, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, 2007.
- J. Persky, Retrospectives: Pareto’s law, *The Journal of Economic Perspectives* 6 (1992) 181–192.
- M. Raberto, S. Cincotti, Analysis and simulation of a double auction artificial financial market, *Physica A* 355 (2005) 34–45.
- V. Alfi, M. Cristelli, L. Pietronero, A. Zaccaria, Minimal agent based model for financial markets II: statistical properties of the linear and multiplicative dynamics, *The European Physical Journal B* 67 (3) (2009b) 399–417.
- V. Alfi, M. Cristelli, L. Pietronero, A. Zaccaria, Mechanisms of self-organization and finite size effects in a minimal agent based model, *J. Stat. Mech* P03016.
- J. Yang, The efficiency of an artificial double auction stock market with neural learning agents, *Evolutionary Computation in Economics and Finance* (2002) 85–106.
- P. Jefferies, M. Hart, P. Hui, N. Johnson, From market games to real-world markets, *The European Physical Journal B* 20 (2000) 493–501.
- B. LeBaron, W. Arthur, R. Palmer, Time series properties of an artificial stock market, *Journal of Economic Dynamics & Control* 23 (1999) 1487–1516.
- D. Gode, S. Sunder, Lower bounds for efficiency of surplus extraction in double auctions, in: D. Friedman, J. Rust (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*. Santa Fe Institute Studies in the Sciences of Complexity, Perseus Publishing, Cambridge, 199–219, 1993.
- M. Aloud, M. Fasli, Exploring Trading Strategies and their Effects in the FX Market, *Computational Intelligence*, doi: 10.1111/coin.12085. .
- A. Atkinson, A. Harrison, *Distribution of Total Wealth in Britain*, Cambridge University Press, Cambridge, 1978.
- K. Arrow, The changing face of economics, in: D. Colander, R. Holt, J. Rosser (Eds.), *Conversations with cutting edge economists*, University of Michigan Press, Ann Arbor, 2004.

Appendix

A Systematic Exploration of the Origin of the FX Market Stylized Facts

This section presents the full obtained results of the different experiments reported in Section 6.

Number of Agents

The FX market involves a large number of participants and we suppose that the total number of agents N in an ABM may be critical in reproducing the trading activity in the FX markets. To examine the impact of the number of agents on the dynamics of the market behaviour, we performed several experiments in which we scaled N from 10^2 to 10^5 .

Seasonality results: Figure 3 shows the intraday seasonality of trade numbers from the simulation results for different values of N . As can be seen from Figure 3, for the different values of N in the market, the dynamics of intraday seasonality of trade numbers resemble those observed in the empirical data. Conversely, as per Figure 4, the dynamic of intraday seasonality of trade volumes for small values of N ($N = 10^2$) is contrary to the observations from the empirical data. In the case of $N = 10^2$ (Figure 4 - a), we can see that there is a degree of continuous fluctuation in the intraday seasonality of trade volumes which is not the case in reality. In the case of $N = 10^5$ (Figure 4 - d), the dynamic of intraday seasonality of trade volumes moves within a narrow level of the resistance, whereas in reality it is wider (Figure 1). Table 6 shows for the intraday seasonality statistics for different values of N , the t-statistic values along with the Pearson correlation values where the t-critical value two-tail = 2.07. According to paired t-tests with regard to the hypothesis ($\mu_R - \mu_S = 0$), the means of the intraday seasonality statistics of trading activity from the simulation results with different values of N , are roughly equivalent as those for the real FX market, with a 95% confidence level. Nevertheless, the trade volumes correlation between the seasonality statistics of real data and simulation data with a fewer agents ($N = 10^2$) is low.

Correlation results: Table 7 reports the correlation coefficients of the different market trading activities from the simulation results with different values of N . The number of agents N does not affect the correlation between (a) number of buy and sell executed orders and (b) number of opening and closing positions. However, such a parameter significantly affects the correlation in terms of the trading volumes in the market. It is clear therefore that the correlation of price volatility and market trading activity depends on a finite value of the total number of agents N , which can be understood in terms of the absence of correlations for smaller values of N (e.g. $N = 10^2$ and $N = 10^3$) and for a very large value of N (e.g. $N = 10^5$).

Scaling laws results: The four scaling laws are not exhibited in the simulation results for small values of N (e.g. $N = 10^2$ and $N = 10^3$) and for a very large value of N (e.g. $N = 10^5$) (as shown from Figure 5 and Table 8). Table 8 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. From Table 8, it is clear that the adjusted R^2 values for $N \in \{10^2, 10^3, 10^5\}$ are not close to 1, which implies an insignificant fit. From Figure 5, the simulation results with few agents (e.g. $N = 10^2$ and $N = 10^3$) in the market do not exhibit the quantitative relationships of the trade volumes scaling law. On average, an OS event contains roughly three times as many trade volumes as a DC event, which in contrary to the observations with regard to the empirical data. These results emphasise the importance of the total number of agents N in an ABM.

Overall, the results indicate the sensitivity of the total number of agents N in an ABM in reproducing the trading activity in the high-frequency FX market. One can see that the full set of stylized facts are obtained only for a finite value of $N = 10^4$. This is due to a small N in the market affecting the intensity of trading activity. One might presume that a very large number of agents in the market would lead to the reproduction of the market stylized facts, given that increasing N induces a higher intensity of trading activity in the market. The opposite is, in fact, the case: a very large number of agents affects the equilibrium of demand and supply in the market, which naturally leads to the price moving in one direction, and essentially affecting the appearance of the market stylized facts.

The impact of the total number of agents has been shown for the Kim and Markowitz and the Lux and Marchesi models in (Egenter et al., 1999). In these two models, the stylized facts appear for a certain range of the total number of agents, and in most cases the stylized facts do not appear for a large number of agents. In the Lux and Marchesi model, the total number of agents affects the excess demand, and therefore affects the




Figure 3: The intraday seasonality of trade numbers from the simulation results for different total number of agents N in the ABFXM where (a) $N= 10^2$, (b) $N= 10^3$ (c) $N= 10^4$ and (d) $N= 10^5$.

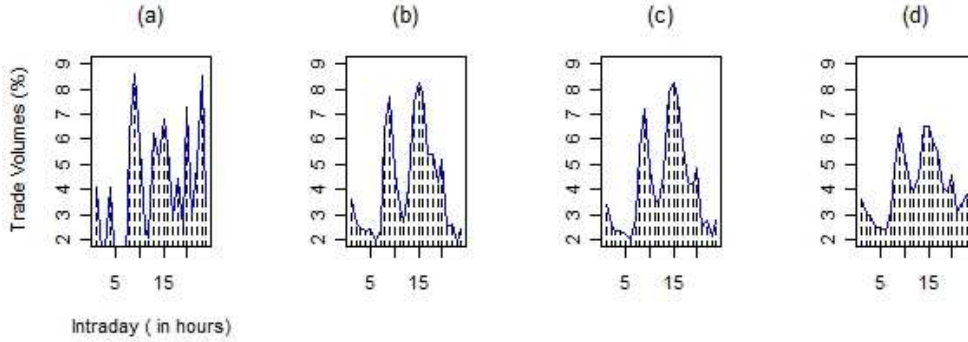


Figure 4: The intraday seasonality of trade volumes from the simulation results for different total number of agents N in the ABFXM where (a) $N= 10^2$, (b) $N= 10^3$ (c) $N= 10^4$ and (d) $N= 10^5$.

appearance of the stylized facts of the price time series. A study of Chen and Yen's Artificial Stock Market shows that the stylized facts tend to vanish when the total number of agents increases (Yeh, 2001).

Initial Activation Condition

In our ABFXM and at the start of the simulation, not all agents are active in the market. We have modelled this as an initial activation condition (INAC) which is, in essence, a probability equal to 10^{-2} . This is to emulate the fact that in real life, not all agents are active all the time. To examine the impact of INAC on the stylized facts, we switch it off and hence all agents are active right from the start. Our results suggest that the incorporation of the INAC is actually quite important for the emergence of stylized facts in the ABFXM.

Seasonality results: A straightforward inspection of Figure 6 shows that the intraday seasonality dynamics are contrary to the observations from the empirical data (Figure 1 (a)). The two peaks of trading activity occur during the early hours of the day in contrast to reality (the peaks occur when the London session is in operation

	(a)		(b)		(c)		(d)	
	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	-1.01E-09	+ 0.98	-7.09E-10	+ 0.98	2.03E-10	+ 0.90	-2.82E-10	+ 0.98
Trade volumes	1.08E-10	+ 0.62	4.64E-16	+ 0.87	1.11E-10	+ 0.86	-9.87E-10	+ 0.97
Opening positions	3.00E-09	+ 0.98	-2.05E-09	+ 0.98	2.02E-10	+ 0.91	5.83E-10	+ 0.98
Closing positions	1.06E-09	+ 0.98	2.27E-09	+ 0.98	-6.08E-10	+ 0.91	5.41E-10	+ 0.95

Table 6: T-statistic (*tsv*) values and Pearson correlation (corr) values where the t-critical value two-tail = 2.07 - Intraday seasonality statistic from simulation results for different total number of agents N in the ABFXM where (a) $N= 10^2$, (b) $N= 10^3$ (c) $N= 10^4$ and (d) $N= 10^5$.




Figure 5: Scaling laws are plotted from simulation results for different total number of agents N in the ABFXM where (a) $N=10^2$, (b) $N=10^3$ (c) $N=10^4$ and (d) $N=10^5$.

	10^2	10^3	10^4	10^5
Trade numbers and volumes	+ 0.33	+ 0.91	+ 0.92	+ 0.89
Numbers of buy and sell executed order	+ 0.99	+ 0.98	+ 0.99	+ 1
Numbers of opening and closing positions	+ 0.99	+ 0.97	+ 0.98	+ 1
Intraday price volatilities and trader numbers	- 0.14	+ 0.01	+ 0.46	- 0.03
Intraday price volatilities and volumes	- 0.04	+ 0.01	+ 0.31	- 0.08

Table 7: Correlation coefficients computed for the different EUR/USD trading activity from the simulation results with different values of number of agents N .

	DC		OS		DC		OS	
	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE
	(a)				(b)			
Trade numbers	0.74	0.63	0.67	0.63	0.72	0.62	0.74	0.56
Trade volumes	0.79	0.58	0.69	0.58	0.71	0.67	0.74	0.55
Opening positions	0.78	0.55	0.68	0.63	0.73	0.57	0.73	0.57
Closing positions	0.79	0.53	0.67	0.62	0.69	0.69	0.74	0.55
	(c)				(d)			
Trade numbers	0.91	0.37	0.90	0.31	0.72	0.58	0.72	0.60
Trade volumes	0.94	0.35	0.94	0.31	0.72	0.63	0.71	0.59
Opening positions	0.92	0.35	0.93	0.31	0.72	0.58	0.72	0.60
Closing positions	0.91	0.39	0.90	0.35	0.72	0.58	0.72	0.60

Table 8: The adjusted R^2 values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from the simulations' results for different total number of agents N where (a) $N= 10^2$, (b) $N= 10^3$ (c) $N= 10^4$ and (d) $N= 10^5$.

(first peak) and the London and New York sessions coincide (second higher peak)). According to the results reported in Table 9 for the paired t-test, the means of the seasonality statistics of the FX market trading activity are roughly equivalent to the simulation results without the inclusion of the INAC. But the correlations between the seasonality statistics of real data and simulation data are extremely low.

Correlation results: As shown in Table 10, the inclusion of the INAC does not affect the correlation between (i) trade numbers and volumes, (ii) number of buy and sell orders and (iii) number of opening and closing positions. However, it does affect the correlation between the intraday price volatilities and trade numbers and volumes. Such negative correlations may be due to the high intensity of trading activity that occurs during the launch of the simulation (as all agents are active).

Scaling laws results: The four scaling laws are not exhibited in the simulation results when the INAC is switched off (as shown from Figure 7 and Table 11). Table 11 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. However, from Table 11, it is noted that the adjusted R^2 values are not close to those observed in the empirical data.

The activation of the agents initial condition in the market appears to be key in reproducing the FX market trading activity. The absence of an initial activation condition would mean that all agents are initially active at the launch of the simulation, which would result in an extremely high level of liquidity in the market at the simulation opening time, compared with the middle and the end of the trading period. To the best of our knowledge, in the literature, none of the existing ABMs imply a constraint with regard to the agents' initial

	t_{sv}	corr
Trade numbers	-2.05E-10	+ 0.04
Trade volumes	4.64E-16	+ 0.01
Opening positions	-2.07E-10	+ 0.04
Closing positions	-2.05E-09	- 0.04

Table 9: T-statistic (t_{sv}) and Pearson correlation values without the INAC where the t-critical value two-tail = 2.07.




Figure 6: Intraday seasonality of trade numbers without the INAC.

	Correlation coefficient
Trade numbers and volumes	+ 0.95
Numbers of buy and sell executed order	+ 0.99
Numbers of opening and closing positions	+ 0.98
Intraday price volatilities and trader numbers	- 0.01
Intraday price volatilities and volumes	- 0.06

Table 10: Correlation coefficients without the INAC.

activation time in the market. In other words, these ABMs assume that all agents are active in the market from the beginning of the simulation.

Limit Orders

The generation of the two types of limit orders in the market - profit-taking and stop-loss limit orders - has a tangible effect on reproducing the stylized facts of FX market trading activity.

Seasonality results: Figure 8 shows the intraday seasonality of trade numbers and volumes from the simulation results, without the two types of limit orders being generated. By a simple examination, we can observe that the dynamics of intraday seasonality are contrary to observations from the empirical data (Figure 1) for the following reasons:

- The intraday seasonality of trading activity in Figure 8 has two peaks, with the first peak being higher than the second peak. In contrast, the empirical data show that the second peak of the market trading activity is higher than the first peak.
- There is a continuous fluctuation at the end of the intraday seasonality of trading activity, between 17:00 and 23:00 GMT, but in reality the FX market trading activities decline sharply.

With regard to the paired t-test results reported in Table 12, the means of the seasonality statistics of the FX

	DC		OS	
	Adj. R ²	SE	Adj. R ²	SE
Trade numbers	0.73	0.59	0.73	0.57
Trade volumes	0.74	0.60	0.72	0.57
Opening positions	0.74	0.56	0.72	0.58
Closing positions	0.72	0.63	0.73	0.56

Table 11: The adjusted R² values and standard errors (SE) of the scaling laws without the INAC.




Figure 7: Scaling laws are plotted from simulation results without the INAC. The x-axis shows the price moves thresholds of the EUR/USD observations.

	<i>tsv</i>	corr
Trade numbers	-4.06E-16	+ 0.86
Trade volumes	3.71E-10	+ 0.91
Opening positions	-4.06E-10	+ 0.87
Closing positions	2.31E-10	+ 0.77

Table 12: T-statistic (*tsv*) values and Pearson correlation values where the t-critical value two-tail = 2.07 - Intraday seasonality statistic from the simulation results without the two types of limit orders being generated.

market trading activity are roughly equivalent to the simulation results without the generation of limit orders in the market.

Correlation results: Table 13 reports the correlation coefficients of the different market trading activities from the simulation results without the generation of limit orders in the market. The inclusion of the generation of limit orders does not greatly affect the correlation between (a) trade numbers and volumes, (b) the number of buy and sell executed orders and (c) the number of opening and closing positions. However, the generation of limit orders in the market affects the correlation between the intraday price volatilities and trade numbers and volumes.

Scaling laws results: The four scaling laws are not exhibited in the simulation results without the generation of the two types of limit orders (as shown from Figure 9 and Table 14). Table 14 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. From

	Correlation coefficient
Trade numbers and volumes	+ 0.81
Numbers of buy and sell executed order	+ 0.99
Numbers of opening and closing positions	+ 0.72
Intraday price volatilities and trader numbers	- 0.05
Intraday price volatilities and volumes	+ 0.01

Table 13: Correlation coefficients computed for the different EUR/USD trading activity from the simulation results without the two types of limit orders being generated.

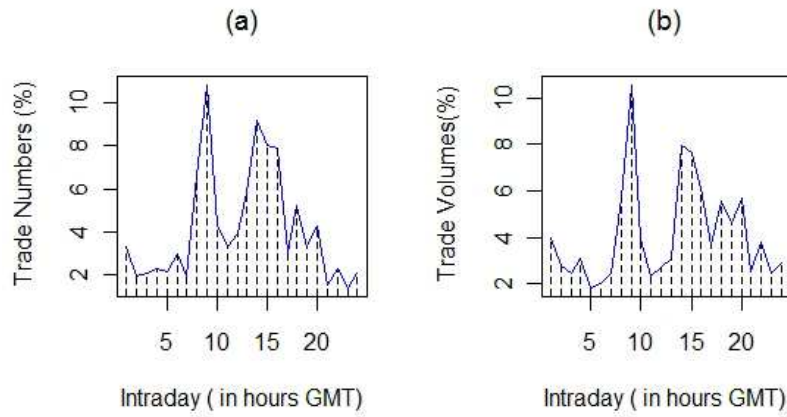


Figure 8: The intraday seasonality of trade (a) numbers and (b) volumes from the simulation results without the two types of limit orders being generated.




Figure 9: Scaling laws are plotted from simulation results without the two types of limit orders being generated. The x-axis shows the price moves thresholds of the EUR/USD observations.

	DC		OS	
	Adj. R^2	SE	Adj. R^2	SE
Trade numbers	0.77	0.55	0.71	0.58
Trade volumes	0.75	0.61	0.70	0.57
Opening positions	0.77	0.55	0.71	0.58
Closing positions	0.95	0.61	0.67	0.64

Table 14: The adjusted R^2 values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from the simulation results without the two types of limit orders being generated.




Figure 10: The intraday seasonality of trade (a) numbers and (b) volumes from the simulation results where all agents are endowed with (1) the same wealth and (2) different amounts of wealth, using a continuous uniform distribution.

Table 14, the adjusted R^2 values are not close to 1 which highlights the importance of the generation of limit orders in the market. From Figure 9, the simulation results without the generation of limit orders do not exhibit quantitative relationships amongst the four scaling laws. On average, an OS event contains roughly four times as many trade numbers and volumes as in a DC event. In addition, an OS event contains about three times as many numbers of opening and closing positions as a DC event. These observations are in contrast to the observations in the empirical data.

The reported results in this section confirm that the generation of limit orders represents an important element in reproducing the trading activity in the FX market. From the results obtained, we can observe that the stylized facts are in contrast to those observed in real FX market data. These unlikely results might be caused by the lack of a comprehensive trading strategy in which the agents should lock a certain profit realization, or seek to limit the potential loss with regard to their open positions. The generation of limit orders is not only responsible for the stylized facts of the intraday dynamics of market trading activity, but limit orders also play a key role in stabilizing trading activity in the market. The importance of the generation of limit orders in an ABM is in line with the results acquired by the studies in (Daniel, 2006; Martinez-Jaramillo and Tsang, 2009; Martinez-Jaramillo, 2007).

Initial Wealth Distribution

Before the launch of the simulation, all agents are endowed with different amounts of wealth based on a power law distribution. The choice of a power law distribution for the agents' initial wealth is capable of reproducing the stylized facts of FX market trading activity.

In order to test the impact of different distribution methods for the agents' initial wealth, we performed some controlled experiments using two different settings. The first setting assumed that all agents are endowed with the same wealth before the simulation starts as has been done by (Daniel, 2006; Martinez-Jaramillo and Tsang, 2009; Levy et al., 1994; LeBaron, 2002; Kim and Markowitz, 1989; Martinez-Jaramillo, 2007). For the second setting, all agents are endowed with different amounts of wealth, using a continuous uniform distribution.

Seasonality results: Figure 10 shows the intraday seasonality of trade numbers and volumes using the two different settings for the initial wealth distribution in the simulation. We can observe by a simple inspection that the dynamics of intraday seasonality are realistic except for the intraday seasonality of trade volumes in the case of a continuous uniform distribution (Figure 10 - (b-2)). We can spot that the first peak is higher than the second which is not the case in reality where the second peak of the market trading activities is higher than the first one. Table 15 shows for the intraday seasonality statistic of the two different settings of the initial wealth distribution, the t-statistic values along with the Pearson correlation values where the t-critical value two-tail = 2.07. According to paired t-tests over the hypothesis ($\mu_R - \mu_S = 0$), the means of the intraday seasonality statistics of trading activity from the simulations results with the two different settings of the initial wealth distribution are roughly equivalent to those for the real FX market, with 95% confidence level.

	(a)		(b)	
	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	-1.06E-09	+ 0.98	-8.07E-09	+ 0.98
Trade volumes	3.38E-16	+ 0.98	4.47E-10	+ 0.98
Opening positions	3.33E-09	+ 0.98	-4.06E-09	+ 0.98
Closing positions	1.44E-09	+ 0.98	-3.09E-09	+ 0.98

Table 15: T-statistic (*tsv*) values and Pearson correlation (corr) values where the t-critical value two-tail = 2.07 - Intraday seasonality statistic from the simulation results where all agents are endowed with (a) the same wealth and (b) different amounts of wealth, using a continuous uniform distribution.

	(a)	(b)
Trade numbers and volumes	+ 0.95	+ 0.90
Numbers of buy and sell executed order	+ 0.99	+ 0.99
Numbers of opening and closing positions	+ 0.99	+ 0.98
Intraday price volatilities and trader numbers	+ 0.01	- 0.01
Intraday price volatilities and volumes	- 0.02	- 0.03

Table 16: Correlation coefficients computed for the EUR/USD trading activity from the simulation results where all agents are endowed with (a) the same wealth and (b) different amounts of wealth, using a continuous uniform distribution.

Correlation results: Table 16 reports the correlation coefficients of the different market trading activities from the simulation results with different settings for the initial wealth distribution. The type of initial wealth distribution does not affect the correlation between the different market trading activities. Nevertheless, such a parameter greatly affects the correlation between intraday price volatility and market trading activity. Consequently, it would appear that the correlation between price volatility and market trading activity depends basically on those agents at the launch of the simulation who possess different wealth by means of a power law distribution.

Scaling laws results: For the two different settings of the initial wealth distribution, the four scaling laws are not exhibited in the simulation results (as shown from Figure 11 and Table 17). Table 17 reports, for the two different settings of initial wealth distribution, the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. From Table 17, it is clear that the adjusted R^2 values are not very close to 1, which implies an insignificant fit. From Figure 11, the simulation results in the case of the distribution of agents' initial wealth is based on a continuous uniform distribution, do not exhibit the quantitative relationships of the trade volumes' scaling law. On average, an OS event contains roughly three times as many trade volumes as a DC event, which is contrary to the observations in the empirical data where, on average, an OS event contains roughly twice as many trade volumes as a DC event.

Assuming that all agents start with the same wealth in the market does not affect the seasonality statistical properties of the different market trading activities. However, it affects the correlation between intraday price volatility and trade numbers and volumes. Also, such an assumption has an impact on the exhibition of scaling laws properties of market trading activities. One of the reasons behind these results is that real

	DC		OS		DC		OS	
	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE
	(a)				(b)			
Trade numbers	0.73	0.58	0.74	0.56	0.73	0.59	0.73	0.56
Trade volumes	0.73	0.59	0.73	0.57	0.73	0.63	0.73	0.56
Opening positions	0.73	0.56	0.73	0.57	0.74	0.56	0.73	0.58
Closing positions	0.72	0.61	0.74	0.55	0.72	0.62	0.73	0.55

Table 17: The adjusted R^2 values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from the simulation results where all agents are endowed with (a) the same wealth and (b) different amounts of wealth, using a continuous uniform distribution.




Figure 11: Scaling laws are plotted from the simulation results where all agents are endowed with (a) the same wealth and (b) different amounts of wealth, using a continuous uniform distribution. The x-axis shows the price moves thresholds of the EUR/USD observations.

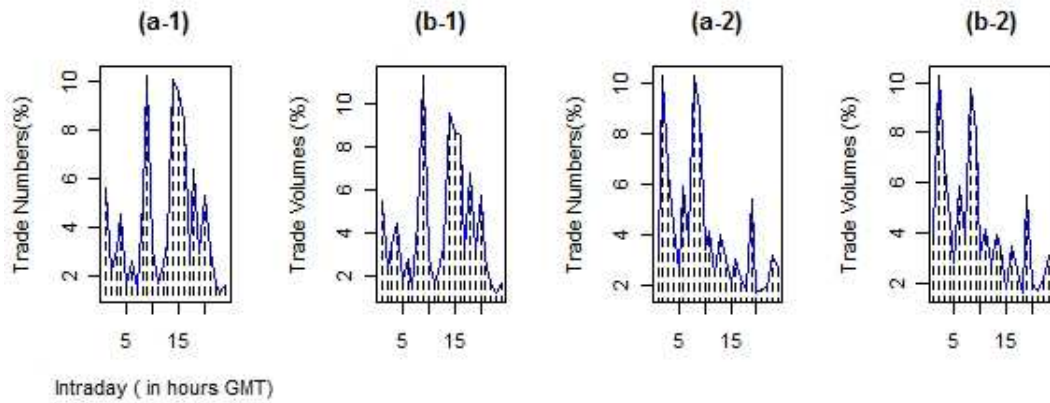


Figure 12: The intraday seasonality of trade (a) numbers and (b) volumes from two different settings of the simulation run for the distribution of the agents' profit objectives: (1) a normal distribution; (2) a power-law distribution.

markets involve heterogeneity in terms of the wealth distribution of the traders. Accordingly, we would argue that the assumption that agents possess the same wealth at the beginning of the market as has been done in (Daniel, 2006; Martinez-Jaramillo and Tsang, 2009; Levy et al., 1994; LeBaron, 2002; Kim and Markowitz, 1989; Martinez-Jaramillo, 2007), is not the way forward in this field.

On the other hand, assuming a continuous uniform distribution for the agents' initial wealth in the market does, to some extent, affect the seasonality statistical properties of the trade volumes. Such an assumption also results in a negative correlation between the intraday price volatility and trade numbers and volumes. Furthermore, it affects the scaling laws behaviour associated with different market trading activities. Asset price and trade volumes evolve endogenously according to the imbalance of demand and supply in the market. This entails an unfairly distributed initial wealth for the agents based on a power-law distribution. From these results we can draw the conclusion that the initial wealth distribution does affect the stylized facts of the market trading activity. The apparent scaling laws observed in real market data find explanations in the presence of uneven distribution in the agents' initial wealth. In the literature, a number of studies found that the income distribution in numerous countries follows a power-law distribution (Pierre, 1965; Steindl, 1965; Atkinson and Harrison, 1978; Persky, 1992). Following on from these studies, we quantify the income distribution in the markets by means of the income distribution of a country's population.

Distribution of the Profit Objective

The precise setting for the agent's profit objective in an ABM is one of the important elements involved in reproducing the stylized facts in the high-frequency FX market. The choice of continuous uniform distribution for the agents' profit objectives rather than other distribution methods, is capable of reproducing, to a certain extent, the stylized facts. In order to investigate the impact of different distribution methods with regard to the agents' profit objectives on the reproduction of the stylized facts, we conducted experiments in which we used a power-law distribution and a normal distribution as alternatives to a continuous uniform distribution.

Seasonality results: Figure 12 shows the intraday seasonality of trade (a) numbers and (b) volumes from two different settings of the simulation run. In the first setting, the agents' profit objectives are based on (1) a normal distribution, while in the second setting (2) on a power-law distribution. We can observe by a simple examination that the dynamics of intraday seasonality are contrary to observations from the empirical data (Figure 1). In the case of a normal distribution (Figure 12 - 1), there are continuous fluctuations in the dynamics of intraday seasonality which is contrary to what is observed in the real data. In the case of the power-law distribution (Figure 12 - 2), it is not possible to see the well-known double U-shape pattern of intraday seasonality of FX market trading activity within the various patterns in the graph. The intensity of the trading activity peaks during the first hours of the day (00:00 am – 08:00 am GMT) while it declines sharply throughout the rest of the day. We conducted a number of experiments by varying the scaling exponent of the power law distribution, but this did not change these unrealistic results.

	(a)		(b)	
	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	-2.03E-10	+ 0.81	-6.05E-11	- 0.08
Trade volumes	2.63E-16	+ 0.83	-2.02E-10	- 0.02
Opening positions	-1.03E-10	+ 0.65	-5.05E-11	- 0.11
Closing positions	2.06E-10	+ 0.82	6.06E-11	- 0.09

Table 18: T-statistic (*tsv*) values and Pearson correlation (corr) values, where the t-critical value two-tail = 2.07 - Intraday seasonality statistic for two different settings of the simulation for the distribution of the agents' profit objectives: (a) a normal distribution; (b) a power-law distribution.

	(a)	(b)
Trade numbers and volumes	+ 0.98	+ 0.98
Numbers of buy and sell executed order	+ 0.57	+ 0.60
Numbers of opening and closing positions	+ 0.25	+ 0.31
Intraday price volatilities and trader numbers	- 0.003	-0.01
Intraday price volatilities and volumes	-0.0004	+ 0.01

Table 19: Correlation coefficients computed for the different EUR/USD trading activity for two different settings of the simulation run for the distribution of the agents' profit objectives: (a) a normal distribution; (b) a power-law distribution.

Table 18 shows the t-statistic values along with the Pearson correlation values, where the t-critical value two-tail = 2.07, for the intraday seasonality statistic for the different settings of the agents' profit objectives distribution. According to paired t-tests with regard to the hypothesis ($\mu_R - \mu_S = 0$), the means of the intraday seasonality statistics of trading activity from the simulations' results, using different settings of the profit objectives distribution, are roughly equivalent to those for the real FX market, with a 95% confidence level. However, in the case of the power law distribution, the correlations between the seasonality statistics of real data and simulation data are negative.

Correlation results: Table 19 reports the correlation coefficients of the different market trading activities from the simulation results with different settings of the agents' profit objective distribution. The distribution method in terms of agents' profit objectives does not affect the correlation between the trade numbers and volumes. But the distribution method affects the correlation between (a) the numbers of buy and sell executed orders, and (b) the numbers of opening and closing positions. In addition, such a distribution method affects the correlation between intraday price volatility and market trading activity considerably.

Scaling laws results: For the different settings of the agents' profit objective distribution, the four scaling laws and the six quantitative relationships amongst them are not exhibited in the simulation results (as shown from Figure 13 and Table 20). Table 20 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. It is clear that the adjusted R^2 values are not very close to 1, which implies an insignificant fit.

The reported results in this section emphasise that the price dynamics in the market are a result of the

	DC		OS		DC		OS	
	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE	Adj. R^2	SE
	(a)				(b)			
Trade numbers	0.80	0.58	0.71	0.58	0.84	0.50	0.69	0.60
Trade volumes	0.80	0.58	0.71	0.58	0.84	0.49	0.69	0.60
Opening positions	0.73	0.70	0.73	0.55	0.80	0.57	0.69	0.59
Closing positions	0.94	0.29	0.69	0.61	0.98	0.15	0.69	0.62

Table 20: The adjusted R^2 values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from two different settings of the simulation run for the distribution method for the agents' profit objectives: (a) a normal distribution; (b) a power-law distribution.

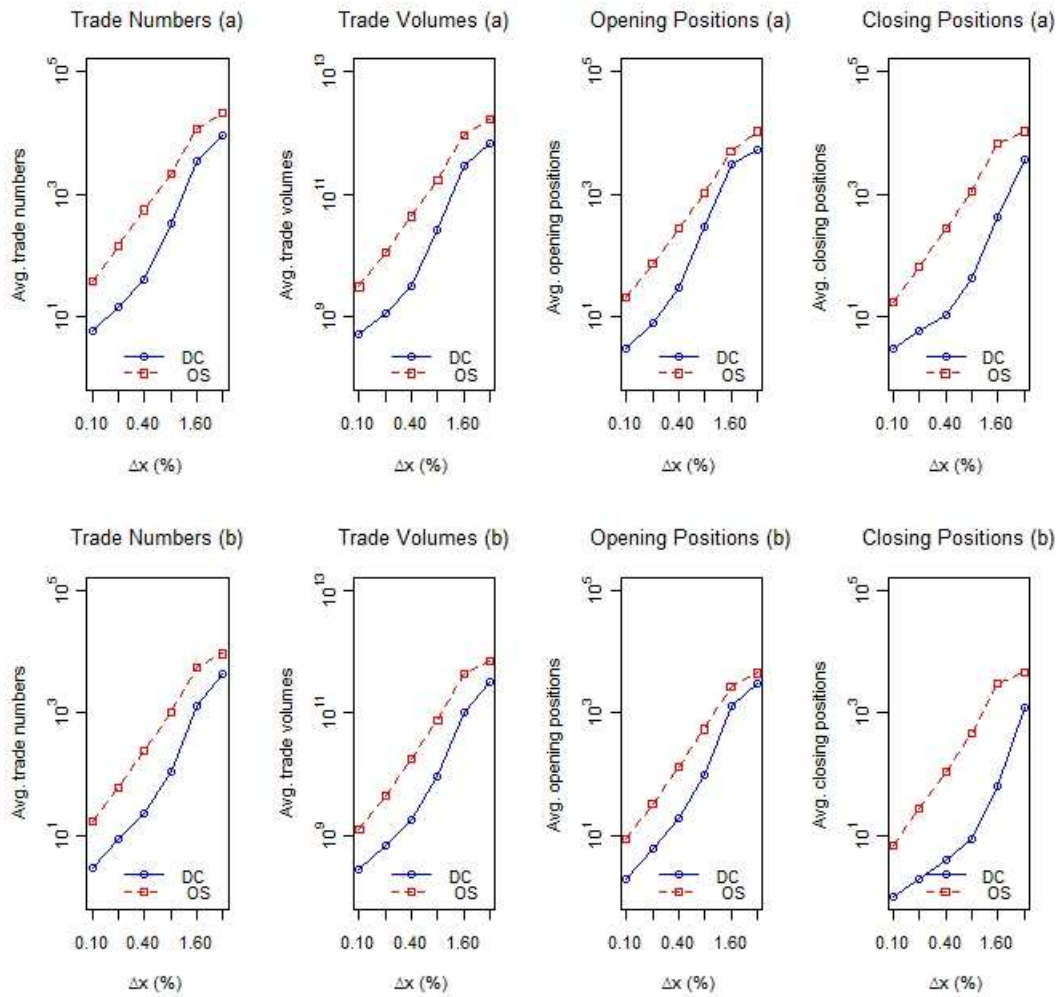


Figure 13: Scaling laws are plotted from two different settings of the simulation run for the distribution method for the agents' profit objectives: (a) a normal distribution; (b) a power-law distribution. The x-axis shows the price moves thresholds of the EUR/USD observations.

existence of heterogeneous expectations amongst traders. According to (Arrow, 2004) “One of the things that microeconomics teaches you is that individuals are not alike. There is heterogeneity, and probably the most important heterogeneity here is heterogeneity of expectations. If we didn’t have heterogeneity, there would be no trade. But developing an analytic model with heterogeneous agents is difficult.” (p. 301).

The power-law and the normal distribution characterize inadequately distributed profit objective values. Consequently, this results in a large group of agents having similar behaviour patterns and preferences whereas, in reality, traders have a variety of preferences and demands. For this reason, choosing a continuous uniform distribution is a good option, given that it provides an equal distribution with a variety of profit objective values for the trading agents.

Risk Appetites

The risk appetite parameter is used to model the cautiousness and aggressiveness of agents with regard to trading. From the microscopic analysis of the OANDA individual traders’ historical transactions, we found on average, that the traders’ risk appetites are higher than their profit objectives (as explained in Section 4.1.3). Accordingly, in our ABFXM, we defined the agents’ risk appetites to be four times that of their profit objectives.

To test the impact of such an element on reproducing the stylized facts, we varied the risk appetites’ parameter in the simulation. Although we have conducted a number of experiments, we report the results of five indicative cases in which the agents’ risk appetites relate to their profit objectives as follows:

1. Extremely high risk: the agents’ risk appetite is higher than four times their profit objectives.
2. High-risk: the agent’s risk appetite is defined as being four times that of its profit objective. The high risk appetite setting is the default value used in our ABFXM.
3. Low-risk: the agent’s risk appetite is equal to the agent’s profit objective.
4. Zero-risk: an agent will close its open position once the unrealized losses for the position fall below zero, which means that the position is in loss.
5. Mixed: prior to the launch of the simulation, the agents are assigned randomly a risk appetite among the three risk appetites 2-4 as described above.

Seasonality results: Figure 14 shows the intraday seasonality of trade numbers and volumes from the different simulation outcomes as per the cases above. When the agents’ trading strategy involves extremely high risk and high risk, the patterns of intraday seasonality of trade numbers are exhibited (high risk is the default setting of the simulation). Meanwhile, when the agents’ trading strategy involves low risk, zero risk or a combination of the three levels of risk, the patterns of intraday seasonality of trade numbers move within a narrow level of the resistance, whereas in reality it is wider. The level of the agents’ risk appetites does not affect the intraday seasonality of trade volumes. Table 21 shows for the intraday seasonality statistic of different settings of the agents’ risk appetites, the t-statistic values, along with the Pearson correlation values. According to the paired t-tests with regard to the hypothesis ($\mu_R - \mu_S = 0$), the means of the intraday seasonality statistics of trading activity from the simulation outcomes with different risk appetite settings, are roughly equivalent to those of the real FX market, with a 95% confidence level.

Correlation results: Looking at Table 22, the agents’ risk appetites do not affect the correlation between (i) trade numbers and volumes, (ii) the number of buy and sell orders and (iii) the number of opening and closing positions. However, the risk appetite results in a negative correlation between the intraday price volatilities and trade numbers and volumes when the risk is extremely-high/low/zero or there is a mixture of risk appetites in the agent population.

Scaling laws results: The scaling laws are not exhibited when the risk is extremely-high/low/zero or there is a mixture of risk appetites in the agent population. This is also clear from Table 23 as the adjusted R^2 values for the four different settings of the agents’ risk appetites parameter (a, c, d, e) are not close to 1, which means an insignificant fit.

The level of the agent’s risk appetite is considered to be an essential element in an ABM. In our agent-based models of traders, an agent’s risk appetite is articulated in terms of the agent’s profit objective. We found that setting a high level of agents’ risk appetites compared to the agents’ profit objectives means that the

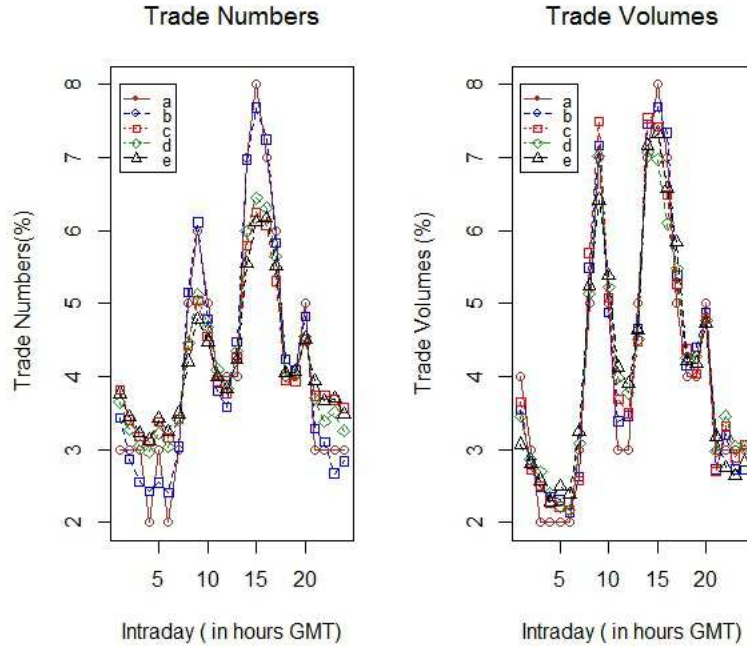


Figure 14: The intraday seasonality of trade numbers and volumes from four different simulation results, where the agents' risk appetites involve (a) extremely high risk, (b) high-risk, (c) low-risk, (d) zero-risk or (e) a combination of the three levels of risk.

	(a)		(b)		(c)	
	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	2.57E-10	+ 0.87	2.03E-10	+ 0.90	-5.07E-10	+ 0.98
Trade volumes	1.18E-10	+ 0.84	1.11E-10	+ 0.86	-1.01E-15	+ 0.98
Opening positions	2.82E-10	+ 0.89	2.02E-10	+ 0.91	-8.05E-10	+ 0.98
Closing positions	-6.44E-10	+ 0.91	-6.08E-10	+ 0.91	8.48E-10	+ 0.98
	(d)		(e)			
	<i>tsv</i>	corr	<i>tsv</i>	corr		
Trade numbers						
Trade volumes	-1.07E-09	+ 0.98	-7.82E-10	+ 0.97		
Opening positions	5.15E-10	+ 0.98	-9.99E-10	+ 0.98		
Closing positions	7.21E-10	+ 0.99	-8.08E-10	+ 0.98		

Table 21: T-statistic (*tsv*) values and Pearson correlation (corr) values, where the t-critical value two-tail = 2.07 - Intraday seasonality statistic from four different simulation results, in which the agents' risk appetites involve (a) extremely high risk, (b) high-risk, (c) low-risk, (d) zero-risk or (e) a combination of the three levels of risk.

	(a)	(b)	(c)	(d)	(e)
Trade numbers and volumes	+ 0.88	+ 0.92	+ 0.87	+ 0.89	+ 0.90
Numbers of buy and sell executed order	+ 0.97	+ 0.99	+ 0.99	+ 0.99	+ 0.99
Numbers of opening and closing positions	+ 0.98	+ 0.98	+ 0.99	+ 0.99	+ 0.99
Intraday price volatilities and trader numbers	- 0.02	+ 0.46	- 0.02	- 0.04	- 0.02
Intraday price volatilities and volumes	- 0.03	+ 0.31	- 0.05	- 0.06	- 0.05

Table 22: Correlation coefficients computed for the different EUR/USD trading activity from four different simulation results, in which the agents' risk appetites involve (a) extremely high risk, (b) high-risk, (c) low-risk, (d) zero-risk or (e) a combination of the three levels of risk.

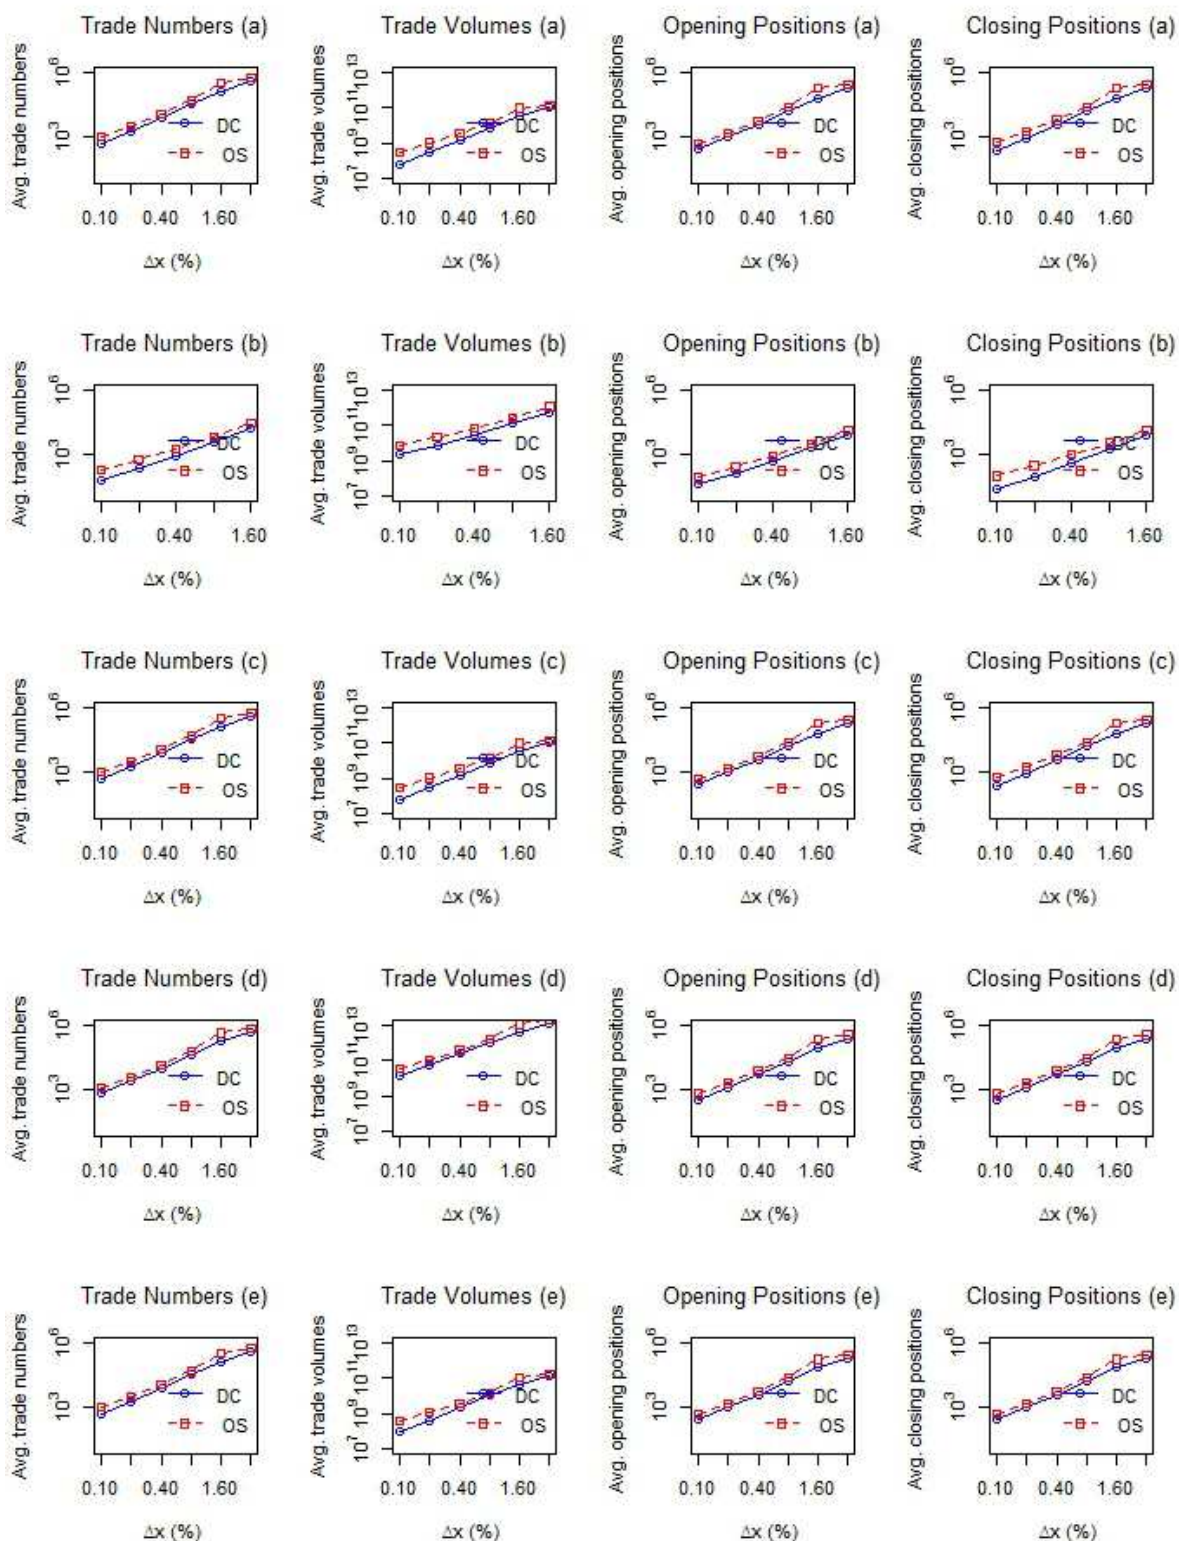


Figure 15: Scaling laws are plotted from four different settings of the simulation run in which the agents' risk appetites involve (a) extremely high risk, (b) high-risk, (c) low-risk, (d) zero-risk or (e) a combination of the three levels of risk. The x-axis shows the price moves thresholds of the EUR/USD observations.

	DC		OS		DC		OS	
	Adj. R ²	SE	Adj. R ²	SE	Adj. R ²	SE	Adj. R ²	SE
	(a)				(b)			
Trade numbers	0.73	0.56	0.73	0.59	0.91	0.37	0.90	0.31
Trade volumes	0.71	0.62	0.70	0.56	0.94	0.35	0.94	0.31
Opening positions	0.73	0.55	0.71	0.62	0.92	0.35	0.93	0.31
Closing positions	0.74	0.56	0.71	0.60	0.91	0.39	0.90	0.35
	(c)				(d)			
Trade numbers	0.74	0.56	0.73	0.58	0.73	0.57	0.73	0.58
Trade volumes	0.72	0.63	0.72	0.56	0.72	0.66	0.72	0.55
Opening positions	0.74	0.56	0.73	0.60	0.74	0.56	0.73	0.58
Closing positions	0.74	0.56	0.73	0.58	0.73	0.59	0.73	0.60
	(e)							
Trade numbers	0.74	0.56	0.73	0.58				
Trade volumes	0.73	0.58	0.73	0.57				
Opening positions	0.74	0.56	0.73	0.58				
Closing positions	0.74	0.56	0.73	0.58				

Table 23: The adjusted R² values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from four different settings of the simulation run in which the agents' risk appetites involve (a) extremely high risk, (b) high-risk, (c) low-risk, (d) zero-risk or (e) a combination of the three levels of risk. The x-axis shows the price moves thresholds of the EUR/USD observations.

ABFXM is capable of reproducing, to a certain extent, the trading activity in the market. However, having the agents' risk appetite higher than four times their profit objectives, results in negative correlations between the intraday price volatilities and trade numbers (- 0.02) and volumes (- 0.03). This is because of the implications of high risk activity on an agent's wealth performance, which triggers margin calls. The right setting of the agents' risk appetites parameter significantly affects the correlation between intraday price volatility and market trading activities. In particular, using a low risk appetite results in a negative correlation between intraday price volatility and market trading activities. Such a negative correlation may be affected by the high intensity of trading activities when the price collapses. The other risk settings also did not seem to be able to reproduce the stylized facts of the FX market. Accordingly, the correlation results confirm what has been observed in market data: traders in the high-frequency market tend to hold losing positions for a long period of time.

Contrarian and Trend-Following Strategies

In our ABFXM, we consider a population whose trading strategies are organized into two major classes with equal number of agents: trend-following and contrarian ZI-DCT0 trading strategies. We aim to study the impact of using just one of these strategies for the entire population and hence we conducted experiments where the adopted trading strategy in the market was either the trend-following or the contrarian strategy.

Also, we examined the effect of the existence of different percentages of contrarian agents in comparison to trend-following agents on the exhibition of the stylized facts. Based on the results obtained, the examination can be divided into two cases, with each case involving a series of experiments. In the first case, the number of contrarian agents is between 30% - 70% of the total population. In the second case, the number of contrarian agents is more than (i) 70% or (ii) less than 30% of the total population: these two sub cases generate very similar results due to a very large number of agents having the same trading strategy. Hence, we only present an example of the latter in this section due to the similarity of the results and for brevity.

In addition, a series of experiments were conducted to study the impact of the inclusion of random trades in the market; random trades have been incorporated in the ABFXM to emulate the fact that in reality traders take decisions influenced by exogeneous to the market events.

Seasonality results: Figure 16 shows the intraday seasonality of trade numbers from the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive of

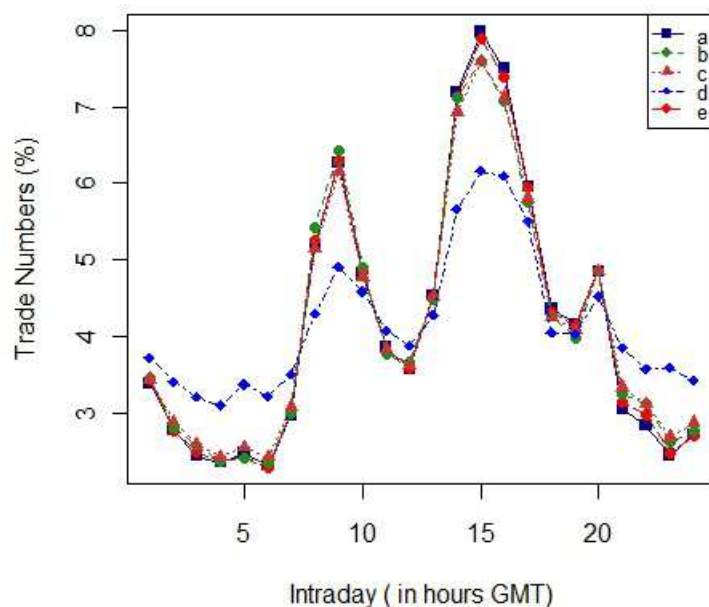


Figure 16: The intraday seasonality of trade numbers where the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive random trades.

random trades from the agents' trading strategy. The dynamics of intraday seasonality of trade numbers resemble those observed in the empirical data where the double U-shape pattern is exhibited. However, in the case (d) the ABFXM populated with a mixture of a mixture of 70% contrarian and 30% trend-following agents, the patterns of intraday seasonality move within a narrow level of the resistance, while in the real data they are wider. This may possibly be due to the imbalance of orders generated by the agents in the market. According to the results reported in Table 24 for the paired t-test, the means of the intraday seasonality statistics of trading activity from the simulation outcomes, are roughly equivalent to those of the real FX market, with a 95% confidence level.

Correlation results: From Table 25, we can observe that the adoption of single strategies does not affect the correlation between trade numbers and volumes, number of buy and sell orders and number of opening and closing positions. However, it results in a negative correlation between the intraday price volatilities and trade numbers and volumes. This implies that the price behaviour in the market is the result of a balance of various trading strategies and levels of activity. The same is the case when random trades are excluded: although there is no impact on the correlation between various trading quantities, the intraday price volatilities and trade numbers/volumes are affected.

Scaling laws results: When either of the strategies is adopted by the whole population and when random trades are excluded, then the scaling laws are not exhibited in the simulation results. The results in Table 26 show that the adjusted R^2 values are not close to 1, which entails an insignificant fit. Also, the estimated standard errors are higher than the estimated errors for scaling law relationships observed in the empirical data. The results of case (c) the ABFXM populated with a mixture of 40% contrarian and 60% trend-following agents, exhibit, the full set of the FX market trading activity stylized facts.

In our ABFXM, we limit the trading strategies to two classes of trend-following and contrarian trading strategies, in which agents can enter or leave the market, depending on the price behaviour. This is an extreme simplification with regards to modelling the market traders' behaviour. The results from our ABFXM show that the simplification in the two classes of trading strategies is enough to achieve an understanding, to a certain extent, of the original market trading activity stylized facts. However, the possibility of a wider class of trading strategies in an ABFXM, and the traders' decisions on the basis of previous performance, is without doubt

	(a)		(b)		(c)	
	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	1.36E-09	+ 0.99	-1.76E-09	+ 0.99	-4.92E-09	+ 0.99
Trade volumes	4.32E-05	+ 0.99	-6.06E-05	+ 0.99	5.67E-10	+ 0.99
Opening positions	1.35E-09	+ 0.99	-1.70E-09	+ 0.99	-4.92E-09	+ 0.99
Closing positions	1.34E-09	+ 0.99	-1.72E-09	+ 0.99	-4.90E-09	+ 0.99
	(d)		(e)			
	<i>tsv</i>	corr	<i>tsv</i>	corr		
Trade numbers	-1.23E-16	+ 0.98	1.74E-09	+ 0.99		
Trade volumes	7.20E-10	+ 0.98	-1.59E-05	+ 0.99		
Opening positions	-1.21E-16	+ 0.99	1.72E-09	+ 0.99		
Closing positions	-1.21E-16	+ 0.98	1.72E-09	+ 0.99		

Table 24: T-statistic (*tsv*) and Pearson correlation (corr) values, where the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive random trades.

	(a)	(b)	(c)	(d)	(e)
Trade numbers and volumes	+ 0.92	+ 0.94	+ 0.95	+ 0.88	+ 0.91
Numbers of buy and sell order	+ 0.96	+ 0.99	+ 0.99	+ 1	+ 0.99
Numbers of opening and closing positions	+ 0.96	+ 0.98	+ 0.98	+ 1	+ 0.98
Intraday price volatilities and trader numbers	- 0.01	- 0.03	+ 0.39	- 0.03	- 0.01
Intraday price volatilities and volumes	- 0.03	- 0.02	+ 0.26	- 0.10	- 0.05

Table 25: Correlation coefficients computed for the EUR/USD trading activity where the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive random trades.

	DC		OS		DC		OS		DC		OS	
	R ²	SE	R ²	SE	R ²	SE	R ²	SE	R ²	SE	R ²	SE
	(a)				(b)				(c)			
Trade numbers	0.73	0.59	0.73	0.58	0.73	0.59	0.73	0.57	0.97	0.35	0.95	0.36
Trade volumes	0.73	0.62	0.73	0.56	0.74	0.60	0.72	0.57	0.97	0.35	0.95	0.36
Opening positions	0.74	0.56	0.73	0.58	0.74	0.56	0.74	0.58	0.97	0.35	0.95	0.37
Closing positions	0.72	0.62	0.73	0.56	0.72	0.62	0.73	0.56	0.97	0.36	0.95	0.36
	DC		OS		DC		OS		DC		OS	
	R ²	SE	R ²	SE	R ²	SE	R ²	SE	R ²	SE	R ²	SE
	(d)				(e)							
Trade numbers	0.73	0.59	0.73	0.57	0.73	0.59	0.73	0.56				
Trade volumes	0.74	0.61	0.72	0.57	0.74	0.62	0.72	0.57				
Opening positions	0.73	0.59	0.73	0.59	0.74	0.56	0.73	0.58				
Closing positions	0.73	0.61	0.73	0.58	0.71	0.63	0.74	0.55				

Table 26: The adjusted R² values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events where the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive random trades.

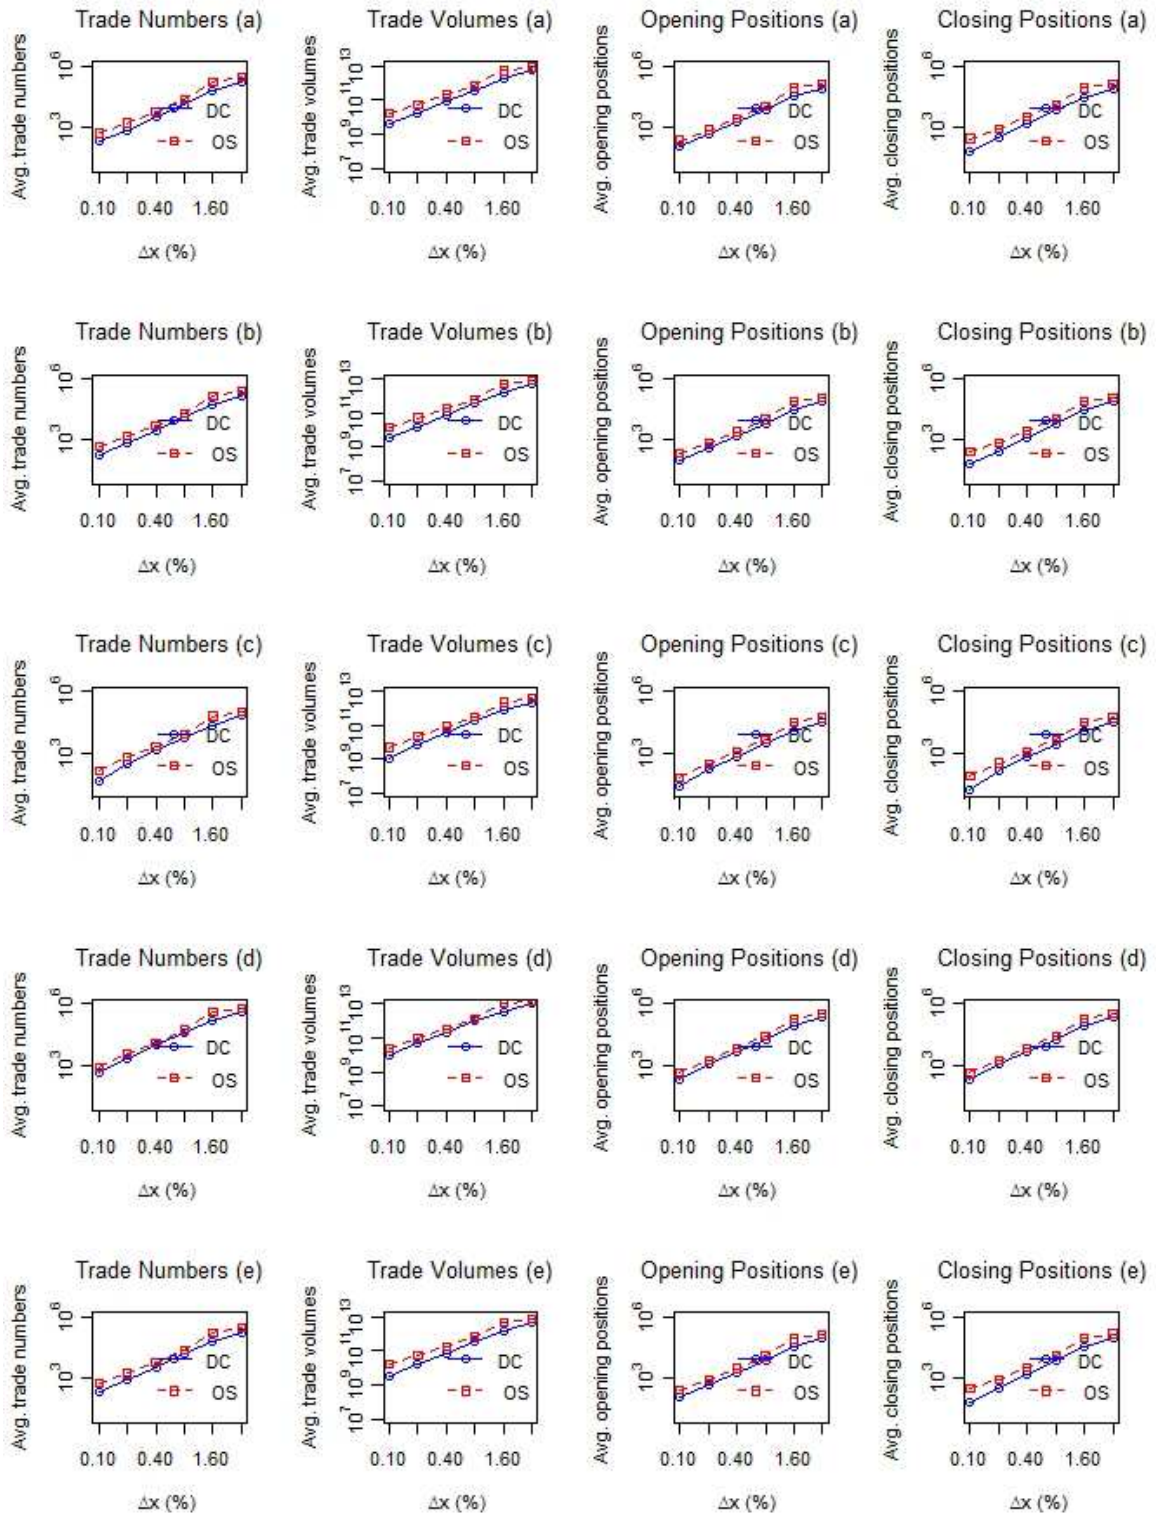


Figure 17: Scaling laws are plotted where the ABFXM populated with (a) just contrarian agents, (b) just trend-following agents, (c) a mixture of 40% contrarian and 60% trend-following agents, (d) a mixture of 70% contrarian and 30% trend-following agents and (e) exclusive random trades. The x-axis shows the price moves thresholds of the EUR/USD observations.

	<i>tsv</i>	corr
Trade numbers	-3.02E-10	+ 0.05
Trade volumes	4.64E-10	+ 0.02
Number of opening positions	-3.05E-10	+ 0.05
Number of closing positions	-3.05E-09	+ 0.04

Table 27: T-statistic (*tsv*) values and Pearson correlation (corr) values, where the t-critical value two-tail = 2.07 - Intraday seasonality statistic from simulation results without incorporating the role of FX market trading sessions.

an interesting area with regard to future exploration. The problem, in this respect, is that the analysis and exploration of trading activities in the ABFXM can become highly misleading and difficult to study, because the broader class of trading strategies may produce a source of complexity for the market analysis.

Therefore, the results support that competition between trend-follower agents and contrarian agents would have a tendency to stabilize the market. The key role played by the coexistence of such agents' trading strategies is one of the main elements responsible for exhibiting stylized facts in an ABFXM. Random trading governs the likelihood of the agents changing their trading strategy, and it is a vital element of the model. In any case, the intraday seasonality statistics of the different market trading activities are exhibited in the limit where the agents' population is composed only of trend-followers or only by contrarian traders, even without random trades taking place. Therefore, these simple experiments show us which stylized facts depend critically on the interaction between different trading strategies. The generalized scaling law properties and the correlation behaviour results appear to be sensitive to the heterogeneity of agent trading strategies.

Similar studies by (Alfi et al., 2009a,b,c) showed that the stylized facts are exhibited in an ABM by introducing the possibility of the agents to enter or leave the market depending on the behaviour of the price. The mechanism to enter or leave the market is linked to the existence of a threshold that the agents use to decide to operate in the market by comparing the threshold with the price movement.

Trading Time Window

Although the FX market operates 24 hours a day and does not have fixed trading time windows as the stock market, the worldwide FX market centres' business hours have a significant effect on characterizing the daily flow of the market's activity. The intensity of trading activity increases during the opening hours of the FX market centres in the morning, whereas it decreases during the closing hours. The overlapping business hours of two or more FX market centres witnesses a high level of trading activity compared to the other times.

Initially, we decided to model the agents without incorporating the role of FX market trading sessions and holidays. The motivation behind our decision was to account for the significant impact of algorithmic trading. Algorithmic trading, also referred to as automated trading, uses computer programs to place orders automatically which may be outside a trader's normal business hours. It has attracted many FX traders and, most importantly, is one of the reasons for the growth of the FX market. However, not modelling business hours has a significant impact on the emergence stylized facts.

Seasonality results: Figure 18 shows the intraday seasonality of trade numbers from the simulation results, without incorporating the role of FX market trading sessions. By a simple examination, we can clearly see that the intraday seasonality of FX market trading activity is not exhibited as it is not possible to spot the double U-shape pattern. This is for the following reasons: the highest volume of trading activity takes place in the first hour of the day, as well as around 08:00 GMT. The trading activity declines sharply from around 09:00 GMT and continues throughout the day, which does not reflect reality. In reality, the daily trading activity increases between 08:00 – 09:00 GMT mirroring the opening hours of the London trading session, and peaks when London and New York trading sessions overlap around 13:00 – 16:00 GMT. According to the results reported in Table 27 for the paired t-test, the means of the seasonality statistics of the FX market trading activity are roughly the same to the simulation results, without incorporating the role of FX market trading sessions. Nevertheless, the correlations between the seasonality statistics of real data and simulation data are extremely low. The correlation results indicate the significant role of incorporating the FX market trading sessions in an agent-based model of FX trading.

Correlation results: Table 28 reports the correlation coefficients of the different market trading activities

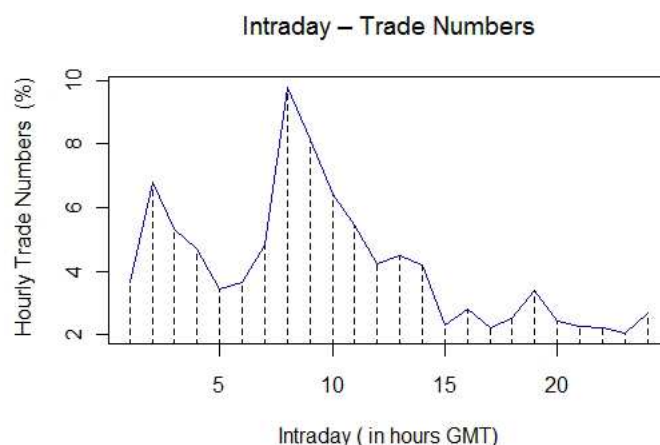


Figure 18: Intraday seasonality of trade numbers from the simulation results, without incorporating the role of FX market trading sessions.

	Correlation coefficient
Trade numbers and volumes	+ 0.87
Numbers of buy and sell executed order	+ 0.99
Numbers of opening and closing positions	+ 0.98
Intraday price volatilities and trader numbers	+ 0.002
Intraday price volatilities and volumes	- 0.09

Table 28: Correlation coefficients computed for the different EUR/USD trading activity from the simulation results without incorporating the role of FX market trading sessions.

from the simulation results, without incorporating the role of FX market trading sessions. The inclusion of the role of FX market trading sessions does not affect the correlation between (a) trade numbers and volumes, (b) numbers of buy and sell executed orders and (c) numbers of opening and closing positions. However, it affects extensively the correlation between the intraday price volatilities and trade numbers and volumes. This means that the time periods of high market trading activity are also the time periods associated with large price volatilities.

Scaling laws results: The four scaling laws measured under DC events are not exhibited in the simulation results when the FX market trading sessions in the agents' trading time window are not incorporated (as shown from Figure 19 and Table 29). In addition, the trade volumes scaling law measured under OS events is not exhibited in the simulation results. Table 29 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. It is clear the adjusted R^2 values are not very close to 1, which implies an insignificant fit.

The results reported in this section stress the importance of incorporating the role of the FX market trading sessions in modelling the agents trading time window. The behaviour of market trading activity exhibits a periodic pattern called seasonal. The seasonal patterns are evidently correlated with the changing presence of

	DC		OS	
	Adj. R^2	SE	Adj. R^2	SE
Trade numbers	0.73	0.62	0.99	0.09
Trade volumes	0.75	0.60	0.72	3.50
Number of opening positions	0.74	0.61	0.99	0.58
Number of closing positions	0.72	0.63	0.99	0.11

Table 29: The adjusted R^2 values of the fits, plus their standard errors (SE), for the scaling laws measured under DC and OS events from the simulation results without incorporating the role of FX market trading sessions.




Figure 19: Scaling laws are plotted from the simulation results without incorporating the role of FX market trading sessions. The x-axis shows the price moves thresholds of the EUR/USD observations.

the main worldwide FX market centres. The results present evidence of a strong correlation between market trading activity and the time of day.

The correlation between market activity and volatility requires us to model a heterogeneous trading time window with the help of information related to the main worldwide FX market centres' business trading hours. The time zones and the typical business hours of the FX market trading sessions (e.g., London, New York, Tokyo) are well known. The results demonstrate the importance of daily seasonal patterns of trading activity in the generation of scaling law properties.

Order Size

In this section, we examine the impact of adopting two approaches that have been used in the literature to determine an agent's order size on the ensuing market transactions and stylized facts. According to the first one, the order size is equal to the agent's wealth (Samanidou et al., 2007). The second approach uses a constant trading proportion parameter that controls the size of an agent's order in terms of the agent's current wealth (LeBaron, 2001; Raberto and Cincotti, 2005; Martinez-Jaramillo and Tsang, 2009). With regard to this second approach, we performed several experiments in which we scaled the trading proportion parameter from 10% to 90%. Due to space limitations and the insignificant changes in reporting all the results, we only report those associated with having the trading proportion parameter equal to 25%, 50% and 75%. For values under 10% the changes were insignificant, while for values above 90%, the results resemble those obtained from using the first approach in which the amount that an agent invests in the market is proportional to the agent's wealth.

Seasonality results: Figure 20 shows the intraday seasonality of trade numbers and volumes from the simulation results with two different settings of the order size parameter: (1) the order size is equal to the agent's wealth, and (2) a constant trading proportion parameter is used to controls the size of an agent's order in terms of the agent's current wealth. As can be seen from Figure 20, the dynamics of the intraday seasonality of trade numbers are realistic for the different settings of the order size parameter. On the other hand, the dynamics of intraday seasonality of trade volumes do not reflect reality where an agent's order size is equal to the agent's wealth. The reason is that the intraday seasonality has two peaks, with the first peak being lower than the second one. In addition, from Figure 20 (Trade volumes), we can clearly see that the trade volume increases suddenly and sharply around 08:00 GMT, which is not the case in reality. Table 30 shows the intraday seasonality statistic of two different settings of the order size parameter, the t-statistic values along with the Pearson correlation values where t-critical value two-tail = 2.07. According to paired t-tests over the hypothesis ($\mu_R - \mu_S = 0$), the means of the intraday seasonality statistics of trading activity from the simulation




Figure 20: The intraday seasonality of trade numbers and volumes from the simulation, results from two different settings of the order size parameter, where: (1) order size is equal to the agent wealth; (2) a constant trading proportion parameter is used to controls the size of an agent’s order in terms of the agent’s current wealth, a constant trading proportion parameter equal to (a) 75%, (b) 50% and (c) 25%.

results with different settings of the initial wealth distribution, are roughly equivalent to the case for the real FX market, with a 95% confidence level.

Correlation results: Table 31 reports the correlation coefficients of the different market trading activities from the simulation results of two different settings of the order size parameter. The type of the order size parameter does not affect the correlation between trade numbers and volumes, numbers of buy and sell executed orders and numbers of opening and closing positions. Nevertheless, it has a significant effect on the correlation between the intraday price volatilities and trade numbers and volumes.

Scaling laws results: The scaling laws are not exhibited in the simulation results with the two different settings of the order size parameter (as shown from Figure 21 and Table 32). Table 32 reports the adjusted R^2 values of the fits, plus their standard errors, for the four scaling laws measured under DC and OS events. We can note that the adjusted R^2 values are not very close to 1, which implies an insignificant fit.

Assuming that a single order size placed by an agent in the market is equal to the agent’s wealth, such as in (Samanidou et al., 2007), does not reproduce the stylized facts. This is because it creates an extremely high trading risk, since an agent invests the full cash it acquires in a single order. This in turn triggers cascaded margin calls, which induce a price fall. A second reason is the lack of heterogeneity in the amount of an agent’s investments during the simulation. The second approach entails that the amount the agent invests in the market will be a fraction of its cash, based on a constant trading proportion parameter (LeBaron, 2001; Martinez-Jaramillo and Tsang, 2009). Similarly to the first one, this second approach induces a lack of heterogeneity of investments in the market. In addition, having a high value in terms of the trading proportion parameter causes cascading margin calls in the market. In contrast, having small values in terms of the trading proportion parameter results in drastic changes in the stylized facts. This is due to the low liquidity in the market and the

	(1)		(2-a)		(2-b)		(2-c)	
	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr	<i>tsv</i>	corr
Trade numbers	- 2.01E-09	+ 0.99	- 9.86E-11	+ 0.99	- 8.54E-11	+ 0.99	- 9.50E-11	+ 0.99
Trade volumes	- 6.06E-16	+ 0.99	- 7.66E-11	+ 0.99	- 7.86E-11	+ 0.99	- 7.59E-11	+ 0.99
Opening positions	- 1.04E-09	+ 0.99	- 9.94E-11	+ 0.99	- 8.64E-11	+ 0.99	- 9.59E-11	+ 0.99
Closing positions	9.59E-10	+ 0.99	1.95E-10	+ 0.99	1.68E-10	+ 0.99	1.88E-10	+ 0.99

Table 30: T-statistic (*tsv*) values and Pearson correlation (corr) values where t-critical value two-tail = 2.07 - Intraday seasonality statistic from the simulation, results from two different settings of the order size parameter, where: (1) order size is equal to the agent wealth; (2) a constant trading proportion parameter is used to controls the size of an agent's order in terms of the agent's current wealth, a constant trading proportion parameter equal to (a) 75%, (b) 50% and (c) 25%.

	(1)	(2 - a)	(2 - b)	(2 - c)
Trade numbers and volumes	+ 0.71	+ 0.81	+ 0.87	+ 0.87
Numbers of buy and sell order	+ 0.99	+ 0.99	+ 0.99	+ 0.99
Numbers of opening and closing positions	+ 0.99	+ 0.99	+ 0.97	+ 0.99
Intraday price volatilities and trader numbers	+ 0.01	+ 0.01	+ 0.01	+ 0.01
Intraday price volatilities and volumes	- 0.06	- 0.02	- 0.04	- 0.02

Table 31: Correlation coefficients computed for the different EUR/USD trading activity from two different settings of the order size parameter, where: (1) order size is equal to the agent wealth; (2) a constant trading proportion parameter is used to controls the size of an agent's order in terms of the agent's current wealth, a constant trading proportion parameter equal to (a) 75%, (b) 50% and (c) 25%.

	DC		OS		DC		OS	
	Adj. R ²	SE	Adj. R ²	SE	Adj. R ²	SE	Adj. R ²	SE
	(1)				(2 - a)			
Trade numbers	0.73	0.58	0.74	0.56	0.76	0.56	0.76	0.56
Trade volumes	0.74	0.61	0.71	0.57	0.74	0.61	0.76	0.56
Opening positions	0.73	0.56	0.73	0.56	0.76	0.56	0.76	0.56
Closing positions	0.72	0.60	0.74	0.56	0.76	0.56	0.76	0.56
	(2 - b)				(2 - c)			
Trade numbers	0.76	0.56	0.76	0.56	0.71	0.56	0.77	0.58
Trade volumes	0.74	0.61	0.76	0.56	0.43	0.70	0.76	0.57
Opening positions	0.77	0.55	0.77	0.55	0.71	0.61	0.76	0.57
Closing positions	0.76	0.57	0.77	0.55	0.71	0.61	0.77	0.56

Table 32: The adjusted R² values and standard errors (SE) of the scaling law fits under DC and OS events from the simulation, results from two different settings of the order size parameter, where an agent's order size is: (1) proportional to its wealth; (2) a fraction of its cash, based on a constant trading proportion parameter equal to (a) 75%, (b) 50% and (c) 25%.

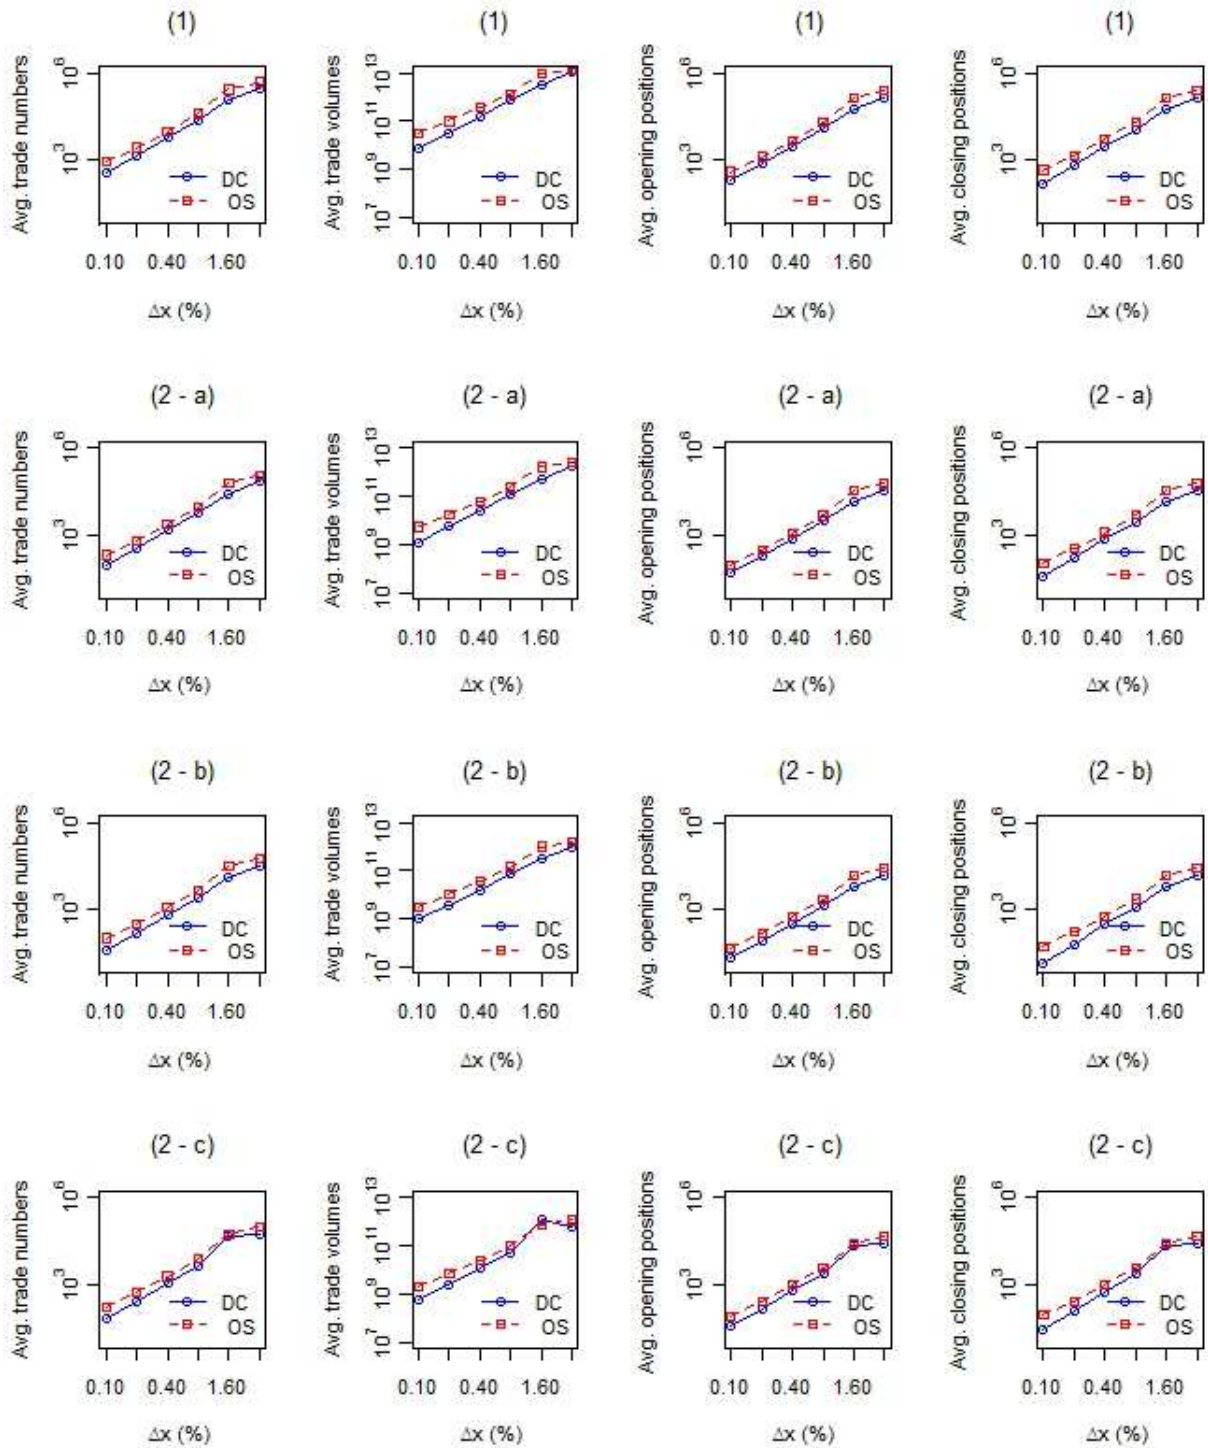


Figure 21: Scaling laws are plotted from the simulation, results from two different settings of the order size parameter, where: (1) order size is equal to the agent wealth; (2) a constant trading proportion parameter is used to controls the size of an agent's order in terms of the agent's current wealth, a constant trading proportion parameter equal to (a) 75%, (b) 50% and (c) 25%.

lack of correlation between the trade numbers and volumes.