Agent-Based Modelling in Directional-Change Intrinsic Time

V. Petrov*†, A. Golub‡ and R. Olsen§

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

We describe an agent-based model where trades happen in event-based time called directional-change intrinsic time. Events are defined as the reversal price moves of a directional-change threshold from a local extreme. The price impact of traded volumes is modelled according to the empirically observed squared root impact function. The time series generated by the agents is characterised by statistical properties typical for foreign exchange rates: low auto-correlation of returns, fat-tailed distribution of returns, aggregated normality, and the price jump scaling law. Furthermore, we introduce and use as a benchmark, the overshoot scaling law, which is an omnipresent feature of liquid markets and relates the expected length of price overshoots to the length of the corresponding directional-change threshold.

Keywords: Agent-based model; Stylised facts; Forex; Directional-change; Intrinsic time; Scaling laws

1. Introduction

According to the Bank for International Settlements, the daily trading volume in Foreign Exchange (FX) markets averaged \$5.1 trillion per day in April 2016[1]. Such substantial volume can have a noticeable impact on the stability of the overall financial system. This volume is generated by the enormous number of transactions made by individual and institutional traders. Traders behaviour is strictly connected to the flow of news related to the financial system (Ederington and Lee 1993, Bauwens *et al.* 2005). Risk management tools capable of foreseeing and efficiently cope with the market impact of any political, environmental or technical change have to be built considering the characteristics and the origin of the market changes initiated by market participants. Due to the large trading volumes and the multitude of trading agents, the FX market is one of the biggest financial systems where agent-based models were extensively applied for its analysis (Aloud *et al.* 2017).

Market participants act on behalf of their perception and interpretation of the available information. Their consolidated behaviour is fully reflected in the prices of the traded financial instrument. This phenomenon is also known as the efficient-market hypothesis (Fama 1970). It encouraged scientists to look into historical time series to study the attributes of the aggregated activity of market participants involved in the trading. Multiple works have been done on the analysis of large amount of historical data in the search for fundamental properties of various financial markets. For example, Bollerslev and Melvin (1994) used more than 300 000 quotes for an empirical analysis of

*Corresponding author. Email: vladimir.petrov@uzh.ch
[1] https://www.bis.org/publ/rpfx16fx.pdf

the bid-ask spread and its relation to the exchange rate uncertainty. Danielsson and de Vries (1997) and Dacorogna *et al.* (2001) used high-frequency data to estimate the fat tails of exchange rate returns. Kozhan and Salmon (2012) employed the dataset of submitted market and limit orders to examine how the information contained in order books could be exploited in simple trading schemes.

A wide range of agent-based models was proposed to explore the observed changes in the markets dynamic. The changes are assumed to be the markets reaction to the actions of individual and aggregated groups of traders inhabiting the market. The models attempt to replicate the evolving behaviour of real market participants by creating artificial agents who impact the market in the same way as it happens in reality. Most of the models represent complex systems, populated by a large number of independent and heterogeneous actors competing with each other (Naciri and Tkiouat 2016). Agent-based models are aimed to reproduce and explain the phenomena of real markets, such as bubbles, crashes, and regime switches (Samanidou *et al.* 2007, Ehrentreich 2007).

There is a wide range of various designs of agent-based models characterised by specific properties intended to imitate phenomenon observed in reality. Despite the diversity of the models, one structural element is always present in their description: the definition of time and how the agents dynamically adapt themselves to its flow. Scientists used to rely on physical time where hours, days or even seasons are selected to measure elapsed periods and intervals between various events (see LeBaron (2006) for the surveys on agent-based models used in finance). Physical time is traditionally employed to describe the interaction between agents and the impact of their trading activity on the market prices. Equally spaced timestamps are assumed to be the points where the actions have to be considered. However, the real market is a complex system with its endogenous non-constant time stream. The activity of that stream is dependent on the inhomogeneous frequency of political, social or environmental events (Guillaume *et al.* 1997). The non-constant volatility typical for the FX market is the main aspect which helps to understand the disadvantages of physical time in agent-based modelling. The volatility of high-frequency markets is the proxy for the market participants activity (Schwert 1989). At the same time, samples of equidistant time intervals of given lengths used to compute volatility can significantly affect its results (Müller *et al.* 1997). Therefore, financial instruments and agent-based models built upon physical time and designed to measure or replicate statistical properties of real time series are naturally limited by the stiffness of selected equidistant intervals.

The limitation of equidistant time periods applied to the financial data analysis can be overcome by applying the concept of directional-change intrinsic time proposed by Guillaume *et al.* (1997). In this intrinsic time representation, events are endogenously defined as reversal price moves measured from local price extremes. All reversals have to be of the size equal to the given directional-change threshold. The proposed measure dissects the real price curve into the sequence of alternating ascending and descending trends. Each elementary trend ends once a new price curve reversal is observed. Continuous price moves in the same direction of the latest reversal are called overshoots. Overshoots represent trend components of the price curve. Only price moves indicating the beginning of the new trend allow the systems clock tick. Thus, this intrinsic time is capable of dealing with the essential properties of price curves such as trends independently of the time intervals chosen to observe them. In other words, intrinsic time is agnostic of the scale and speed of the price evolution registered in physical time.

Trend changes are the primary indicators of the prevailed side of the aggregated behaviour of all market participants (Muth 1961, Cohen *et al.* 2007). The agent-based model, proposed in our work, employs the directional-change intrinsic time framework to monitor trends and dissect the price curve into a collection of directional-change points. These points and price overshoots are used by the agents as signals for their trading activity. Agents, operating in directional-change intrinsic time, judge the market conditions and make decisions on their actions employing prices as the only source of information.

A successful agent-based model should be able to reproduce statistical properties of real financial markets. These properties are mostly known as 'stylised facts' and are copiously discussed in numerous papers (Kaldor 1961, Pagan 1996, Kullmann *et al.* 1999, Gençay *et al.* 2001, Cont 2001, Chakraborti *et al.* 2009). In this work, we test the time series synthetically created by the model against a set of benchmarks. The main target of the test is to confirm that the created time series are characterised by the set of selected statistical properties. The following stylised facts are among the chosen tests: low autocorrelation of returns, fat-tailed distribution of returns, aggregational normality, the price jumps scaling law[1] and the overshoot scaling law. The latter statistical property was chosen in addition to the four stylised facts usually adopted as benchmarks of agent-based models. This scaling law was recently found in a wide range of real high-frequency time series from the FX market and even in arithmetic Brownian motion (Glattfelder *et al.* 2011). It establishes a relation between the average length of overshoot and the corresponding size of the directional-change threshold. The relation states that on average, a directional change is followed by an overshoot of the same magnitude. The sovereignty of the directional-change concept of the flow of physical time makes the overshoot scaling law a very convenient tool for testing the performance of agent-based models. To the extent of our knowledge, it is the first research work where a stylised fact based on the directional-change intrinsic time is used to evaluate the time series generated by a group of heterogeneous agents.

The outline of the remaining paper is as follows. Section 2 illustrates the details of the directional-change intrinsic time framework. It also provides an example of a real price curve dissected into a collection of intrinsic events. Section 3 describes two main components of the agent-based model: the set of artificial agents and the market impact function. In Section 4 the collection of benchmarks used to validate the properties of the generated time series is discussed. In Section 5 we present all obtained results and observed statistical properties of generated time series. In Appendix A the average length of overshoots is derived for the case of Brownian motion with the constant trend. In Appendix B the pseudo-code of the directional-change intrinsic time algorithm is provided. Appendix C concludes the paper with a collection of graphics describing all auxiliary experiments.

2. Intrinsic Time

The number of transactions as well as traded volumes in liquid markets is much lower during holidays or weekends than during working days or after some unexpected but significant news (Chordia *et al.* 2001). This evidence directly contributes to the continuous changes in the financial time series volatility over time (Blattberg and Gonedes 1974, Christie 1982, Scott 1987). Thanks to this non-homogeneous nature of financial markets, one can find drastically distinct price curve evolution within two separate historical time intervals of equal length. One period could be characterised by an instant price drop by several percents and not less instant recovery to the same level (for example, a flash crash) when the other could represent an absolute standstill. The latter often happens when the market is completely inactive due to holidays. Despite this well-known fact, the historical dynamics of financial markets have been mostly analysed using snapshots of market states equally spaced in time. Such snapshots are typically either too infrequent and do not allow capturing of all the available high-frequency information (Zhou 1996) or are too numerous which cause the excessive noisy details of the created set of data (Zhang *et al.* 2005).

Mandelbrot and Taylor (1967) were one of the first researchers to propose an alternative event-based paradigm for modelling and analysing financial time series. Price changes over a fixed number of transactions were studied and compared to the price changes over fixed periods of time. Later, Guillaume *et al.* (1997) extended the early set of alternative data analysis techniques by introducing the directional-change intrinsic time concept. According to that concept, all ticks happen as a result of alternating rising and falling prices (returns) of a certain size. The desired for the observer size of the returns can be arbitrary chosen This allows independent monitoring of the trend changes at various scales. It is especially important taking into account that same price curve can be

[1]Scaling law (power law): the mathematical relationship between two variables that holds true over multiple orders of magnitude.

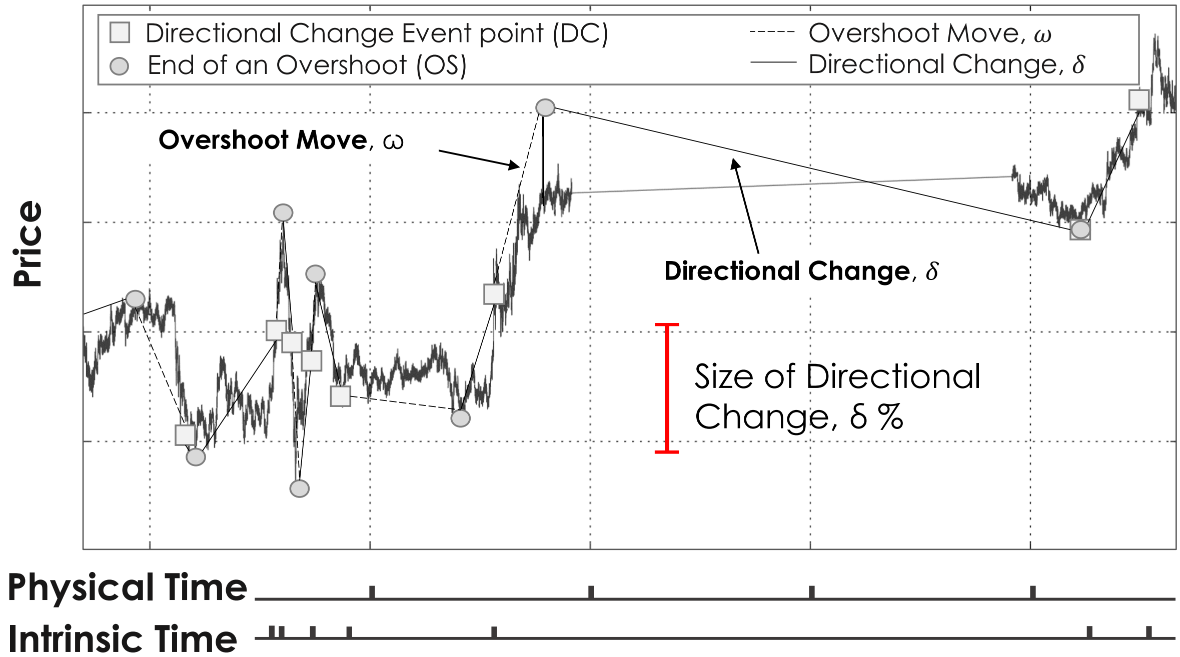

Figure 1. Example of a price curve dissected into a collection of directional changes using a threshold δ. Two types of sections are typical for each consecutive pair of intrinsic events: directional-change (marked by the solid line) and overshoot (dashed line). The direction of the initial mode is chosen to be downward. The first (left) grey square marks the first directional-change event which occurs when the price moves downward by δ percent from the local extreme. The local extreme for the given mode, coincides with the highest observed price (the first grey circle). The mode alternates (upward) as soon as a new directional-change point becomes registered. After this step, the local extreme coincides with the smallest observed price since the time of the latest directional-change. The next upward event is registered when a positive return of the size δ happens measured from the local extreme. After this, the mode alternates again and the dissection process continues.

characterised by different directions of low- and high-scale trends. The phenomenon has been frequently exploited by traders. One specific trading strategy called 'Triple Screen Trading System' was proposed by Elder (2014).

The purpose of the directional-change intrinsic time is to register moments at which the price curve alternates its trend of the given scale. It also finds extreme prices which correspond to local maximum or minimum between two consecutive trend flips. The next paragraph contains a detailed description of the dissection procedure of tick-by-tick data by the directional-change algorithm. An example of a real price curve dissected into a collection of intrinsic events is shown in figure 1.

The latest observed tick price should be taken as the starting point to initialise the dissection of a price curve into a set of directional changes. Then, the relative size of the directional-change thresholds δ and the initial direction of the trend should be selected. There are only two possible states of the trend which correspond to each given threshold δ: mode *up* for the upward or mode *down* for the downward trend. The initial value of the extreme price S_{ext} coincides with the first price used to initialise the algorithm. Each new tick has to be compared to the latest registered extreme. If the current mode is *up* (*down*) and the newest price S_{tick} is higher (lower) than the extreme price S_{ext}, then the S_{ext} takes the value of S_{tick}. Alternatively, the distance between the latest price S_{tick} and the current local extreme S_{ext} should be compared to the size of the threshold δ. If the distance is bigger or equal to δ then the current price is a new directional-change point. At this moment, the mode should be changed to the opposite one and the local extreme reinitialised by $S_{ext} \leftarrow S_{tick}$. The overshoot part ω of the price curve corresponds to the price trajectory between the latest directional-change and the local extreme points.

Overshoot intrinsic events (overshoots events) are registered every time whenever the size of the overshoot ω becomes a multiple of the dissection threshold δ (Golub *et al.* 2017). We employ this category of intrinsic events in addition to the initially proposed in Guillaume *et al.* (1997) directional changes. The overshoot events will be connected to the empirically observed length of the overshoot section with the decision-making mechanism of the agent-based model. As the reader can see from figure 1, the minimum size of overshoot is equal to zero. That can be observed when the price does not follow the trend and makes a reversal right after a new directional-change point. At the same time, there is no the upper bound on the size of the overshoot section. Therefore, there is no a limit on the number of overshoot intrinsic events between two consequent directional changes. The details on the expected length of ω will be provided in Section 4.2 where the overshoot scaling law is described. The pseudo-code of the dissection algorithm is provided in Appendix B.

The recording of the price activity in directional-change intrinsic time does not rely on the exogenous evolution of physical time. Only endogenous price moves define the steps at which a measure should be taken. Interesting properties of the financial time series emerge in this case. Guillaume *et al.* (1997) presented a scaling-law discovered with the help of their new event-based concept. The scaling law establishes the relationship between the chosen thresholds δ and the number of the corresponding trend changes $N(\delta)$ registered within the given time period:

$$N(\delta) = \left(\frac{\delta}{C}\right)^{E}, \tag{1}$$

where $C \in (0, +\infty)$ and $E \in (-\infty, +\infty)$ are scaling law parameters. This scaling law states that the number of directional changes observed in a data sample of the fixed length relates to the size of the chosen threshold as a power law function. It is worth noting that although the size of the threshold δ can range from zero to infinity, the practically selected values depend on the volatility of the analysed process and the frequency of price changes. For example, Glattfelder *et al.* (2011) selected the set of thresholds ranging from 0.01% to 5% to analyse scaling laws of foreign exchange rates. The typical annualised volatility of FX rates is about 20% which means that the selected range covers the majority of price trajectories observed within several years. At the same time, the thresholds are practically bigger applied to the analysis of emerging markets such as their volatility is substantially higher (Petrov *et al.* 2018).

The intrinsic time ticks more often during periods of high volatility and less often when the market is relatively quiet. The latter insures that no important information will be lost as well as no noisy points will enter the created data sample. The visualisation of the statement is provided in figure 1 where the last two directional-change events are registered after a weekend (right part of the plot where the exchange rate does not move). There were no trades therefore no price moves and therefore no information which could enrich the collected recorded. Additionally, the directional-change intrinsic time allows capturing the most valuable moments represented by the local maximums and minimums indicating the end of each trend. All price ticks between the registered intrinsic events are considered as noisy at the scale given by the size of the threshold and become ignored. The prices, neglected by one threshold can at the same time play the role of tipping points at another scale. This multi-scale property of the directional-change intrinsic time can be used to decrease the signal to noise ratio in each specific experimental scenario sensitive to the scale of market moves.

There are various advantages of the event-based approach employed for high-frequency data analysis. They have been successfully applied and described in multiple papers. Among others:

(i) The directional-change algorithm was used to discover numerous scaling laws hidden in the price curves from the FX market. The found properties are expected to improve the inferences we make on the price evolution through analysing liquid markets (Glattfelder *et al.* 2011);

(ii) Several directional-change thresholds of various sizes can be initiated at the same time. The evolution of such multiscale system can be employed to describe prices of the given financial time series. This interpretation of the market dynamic leads to an efficient estimator of the price trajectory unlikeliness. This unlikeliness is later used as the indicator of the current and the forecasted liquidity of the market (Golub *et al.* 2014);

(iii) The number of directional-changes grows fast in periods of high volatility. The number stays small when the volatility is close to zero. Thus, the approach can be used as a volatility estimator of the given time series. The number of events as well as the magnitude of the trend components are the proxy for the volatility size (Petrov *et al.* 2018).

A large number of directional changes can be registered if the size of the threshold δ is significantly smaller than the volatility of the analysed time series. The average numbers of overshoot intrinsic events observed in the upward and downward local trends is approximately equal to each other if the overall trend of the time series is equal to zero. Alternatively, if there is a persistent trend, the price curve tends to overshoot more often when the type of the directional-change mode coincides with the direction of the overall trend. Thus, the average numbers of overshoot events in upward and downward trends will not be equal to each other. However, it is possible to modify the original algorithm in such a way that the trend will have no impact on the balance of the overshoots number. A couple of distinct thresholds can be used to monitor two classes of directional-change modes independently: δ_{up} to register events which happen within the upward trend and δ_{down} for events within the downward trend. The theoretical analysis in Appendix A demonstrates how the trend size affects the expected length of the overshoot sections. As it can be seen from equations (A6) and (A6), the expected size of an overshoot depends on the trend of the market μ and on its volatility σ. For example, if the trend is negative, then the expected upward overshoots $\omega(\delta_{up})$ are equal to downward overshoots $\omega(\delta_{down})$ only if $\delta_{up} < \delta_{down}$. It is important to note that the length of the overshoot section contributes to the number of overshoot intrinsic events. Therefore, the equations from Appendix A can be used as the reference to the expected quantity of intrinsic events per given period of time.

Most of the analytical tools developed in finance were initially built using the assumption that the analysed market does not exhibit any particular trend. However, real price curves are collections of alternating trends of various scales. Therefore, parameters of some trading or risk-management algorithms require adjustments which can compensate nonzero trend typical for the given market and the selected period of time. The directional-change intrinsic time concept, equipped by the aforementioned asymmetric thresholds property, is one of the instruments capable to deal with that trend problem. A couple of directional-change thresholds $(\delta_{up}, \delta_{down})$ can be selected in such a way that the number of upward and downward overshoot intrinsic events will be always equal to each other. However, the price trend cannot be estimated ex-ante. Thus, it is impossible to say which pair of thresholds should be chosen to neutralise the drift expected in the near feature. An efficient solution of the problem was proposed in Golub *et al.* (2017). The authors describe an automated trading algorithm utilising a set of statistical properties accompanying the directional-change intrinsic time. The algorithm efficiently generates profit by capturing the price curve coastline and trading at the moments of expected alternating trends. The equal number of buy and sell trades should happen within periods of the growing and the falling price sections to ensure stable positive returns. In other words, at any moment of time, the algorithm should have precisely tuned pair of upward and downward thresholds which corresponds to the currently present trend. The authors used the accumulated inventory of the trading bot as the proxy for the trend. The selected thresholds should be changed every time when the absolute size of the inventory crosses selected limits. In our work, we are creating a trivial agent-based model where the artificial agents do not keep track of their inventory. Instead, each of them has a unique ex-ante specified pair of non-modified thresholds. Such settings lead to the case when at any moment of time there is an agent whose couple of thresholds captures the current price trajectory in the best possible way (registering equal number of upward and downward overshoot events). The whole set of agents and their parameters will be described in the following section.

3. Structure of the model

Financial markets can be schematically described as the interplay of two components. First, a group of trading agents characterised by the diverse range of behavioural patterns. Second, the exchange rate response on the aggregated agents trades (see, for example, Delage *et al.* (2010)). We stick to this generalised representation to keep our model as simple as possible. Therefore, the agent-based model has: a set of agents buying and selling a fixed volume only at the moment when they register an intrinsic event; the volume impact function, a special algorithm which generates the next price move (return) as a reaction to the aggregated imbalance of the number of buyers and sellers. The latter will be described a few paragraphs below.

The demand and supply create equilibrium in any efficient financial system. In other words, the total buy and sell volumes over given period of time are equal to each other. Nevertheless, our agent-based model does not assume the complete equilibrium. Instead, we postulate that the excess volumes formed by all trading agents can be endlessly bought and sold to and from some external counterparties. The excess volume will be later connected to the price returns through the volume impact function.

The whole life cycle of each agent consists of only two actions: opening a long or short position and waiting for a moment to reverse it. The agents do not exploit each possible opportunity to perform trades (in the same way how it happens in the real world): they flip their positions with probability P_{flip} which makes the performance of the model more realistic. The moments when the agents try to open or flip the positions coincide with the occasion when intrinsic events become observed. The sequence of such moments is unique for each agent since it is determined by the assigned directional-change thresholds δ_{up} and δ_{down}.

For simplicity, all agents trade identical volumes equal to one lot. To flip a position from long to short an agent shorts one unit to close its long position and shorts an additional unit to open a short position. In total two units should be sold. A similar procedure is in place for flipping from a short to a long position, whereas the agent buys two units. Thus, at any iteration there are always N_{long} agents who decided to flip their position from short to long after analysing the latest return and N_{short} agents who decided to become short instead of long. The value $\Delta N_n = N_{long} - N_{short}$ indicates the excess demand or supply at the step n. The variable ΔN_n, also called the net volume, is used by the model to determine the subsequent price change using the volume impact function. A simplified example of a sequence of intrinsic events registered by the trader with parameters $\delta_{up} = 3$ and $\delta_{down} = 2$ is visualised in figure 2.

Real market participants have a diverse set of trading strategies: trading in working days or weekends, using technical or fundamental analysis, preferring high-frequency trading or holding long-term positions (see, for example, the survey of US market provided by Cheung and Chinn (2001)). Good agent-based models aimed to mimic real liquid markets should also be characterised by such varied behaviour. We emulate a complex system by creating a group of intrinsic event agents each of which has a unique set of parameters δ_{up} and δ_{down} which impact their trading patterns. The model does not have a pair of agents with completely the same settings. Each agent interprets a price curve from the point of view of the directional-change intrinsic time. The interpretation depends on the size of the assigned thresholds and their dissimilarity. Some of the simulated traders can register intrinsic events at each new price tick by operating with small enough thresholds (like real high-frequency traders). At the same time, traders with large thresholds register intrinsic time events significantly less frequently which makes their behaviour similar to the one of the long-term investors. Therefore, the selected diverse composition of the model makes the behaviour of individual traders exclusive and the whole range of strategies manifold. As a result, various trading activity patterns are reproduced.

The goal of the agent-based model is to generate a set of returns characterised by the same statistical properties found in the FX market. Each produced return is assumed to be defined in logarithmic terms (log-returns). This assumption makes it possible to directly compare aggregated returns and the size of the directional-change thresholds typically expressed in the percentage

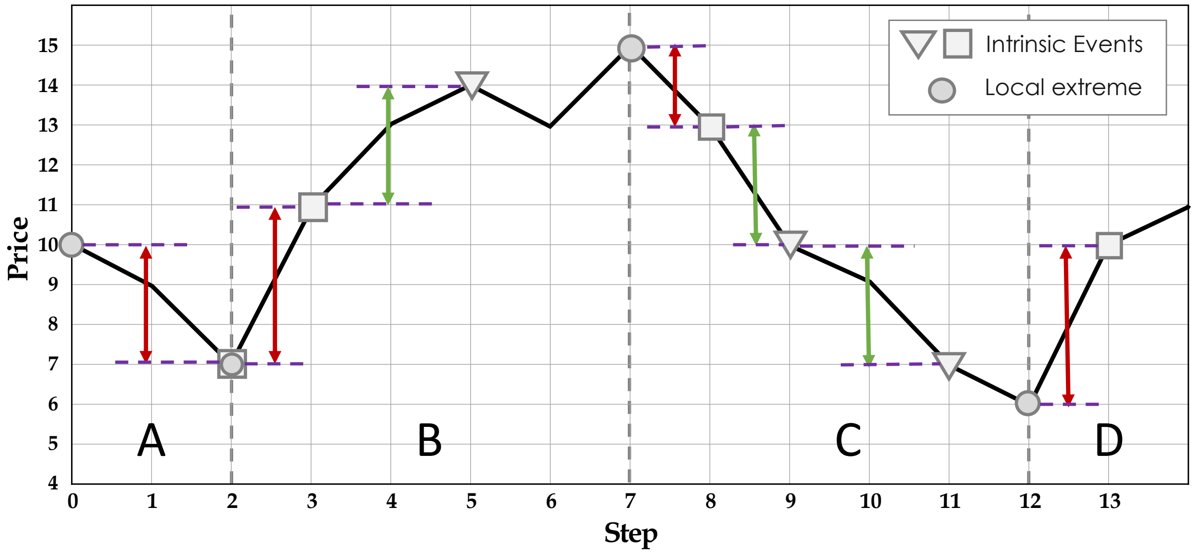

Figure 2. An example of a price curve dissected by a set of intrinsic events at which the agent with parameters $\delta_{up} = 3$ and $\delta_{down} = 2$ opens a position and makes trades. For simplicity, the chosen thresholds are expressed in the absolute price moves instead of the relative ones. The initial price is equal to 10 and the agents initial mode is down. Red arrows mark the distance measured between local extremes and the following directional change points. Green arrows label price moves which lead to the next overshoot intrinsic event. Letters A, B, C and D are put here to tag four intervals of alternating modes. The agent registers its first directional-change intrinsic event as soon as the price goes down from the local extreme by at least two points (labelled by a circle). Since this is the first intrinsic event observed by the agent, it opens its first position at this point (step number 2). Then it is waiting for the next event which happens at step 3 after a big price move by four points up from the latest local extreme (coincides with the previous directional change intrinsic event, a grey square). Independently on its previous trading decision, the agent keeps analysing the price curve. At step 5 the price has made an overshoot move measured from the preceding directional-change point the size of which is equal to up (a grey triangle). This point marks the first overshoot intrinsic event. Though the price continues its overshoot move up, it does not go far enough to trigger another overshoot intrinsic event and the next tipping point becomes a directional change at step 8. The next two overshoot intrinsic events happen at steps 9 and 11. The last directional change concludes the example at step 13.

terms. In other words, returns between two given prices at steps m and n (x_n and x_m) are defined as $r(n, m) \approx \log(x_m) - \log(x_n) = S_m - S_n$ where S_m and S_n are sums of returns accumulated by the steps m and n. Thus, a new directional-change intrinsic event happens when the absolute distance between the latest price and the local extreme becomes bigger or equal to δ. This simplification significantly facilitates all computations and will be used in the rest of the article.

To simplify the model, we represent the whole set of agents as a collection of separate points on a square grid (figure 3). Each point corresponds to a couple of thresholds assigned to one unique trader. Values δ_{down} are located along the horizontal side of the grid (x-axis) and all included in the model thresholds δ_{up} are placed along the vertical side (y-axis). It is worth mentioning that the given setup is built on the assumption that any price return should become a signal for a trade. Thus, the distance between two consecutive values on the same axis was selected to be equal to one price tick. This guarantees that any cumulative return will serve as a new signal for the agent-based model.

The geometrical size of the grid defines the extent to which the agents cover the diversity of various trading patterns. We split the grid by L points horizontally and vertically to assure the symmetry of the trading strategies. Figure 1(a) in Appendix C demonstrates the number of trades performed by all individual agents from the entire grid. As expected, agents with the smallest thresholds located in the left bottom corner make the biggest number of trades (high-frequency traders). The right top corner represents the rarely trading of agents who perform using the biggest thresholds (long-term holders).

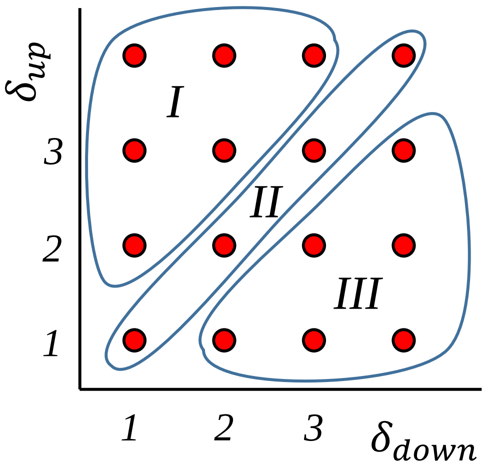

Figure 3. A part of the grid representing the collection of trading agents. Each point corresponds to an agent defined by a set of unique parameters (δ_{up}, δ_{down}). δ_{up} and δ_{down} are the size of upward and downward directional-change thresholds correspondingly. We mark regions with specific properties by numbers I, II and III. In the region I there are traders with the upward directional change thresholds larger than the downward ones ($\delta_{up} > \delta_{down}$). The region II contains 'symmetric' agents ($\delta_{up} = \delta_{down}$). Region III labels all agents with downward thresholds larger than upward ones ($\delta_{up} < \delta_{down}$).).

The grid can visually be divided into three separate sections according to the types of agents activity (figure 3). Traders from section I have upward directional-change thresholds larger than the downward ones ($\delta_{up} > \delta_{down}$). These agents register the equal number of overshoot intrinsic events within local trends upward and downwards only when there is a persistent positive trend. Otherwise, the average number of overshoot intrinsic events registered within the directional-change mode downwards prevail over of the average number within the upward sections. These agents tend to exploit more trading opportunities when the trend is negative or equal to zero. Agents from the region II have equal upward and downward thresholds and are called diagonal agents. The absence of any global trend is the most stable condition for them. In that case, they witness the equal number overshoot intrinsic events within local trends in both directions. Region III marks all agents with upward directional-change thresholds smaller than the downward ones ($\delta_{up} < \delta_{down}$). Their behaviour is the opposite to the behaviour of traders from the region I. For each trader from the region I there is a corresponding opposite agent from the region III. Therefore, the complete set is fully symmetric and thus imitates the counterbalance of different traders in the real financial world. We will show later that one can simulate steady trends of given sizes by deactivating agents from specific parts of the grid.

The lack of demand motivates the supply side to reduce prices and the lack of supply affects the price rises (Walras 2013). This empirical observation was used in our work to define the volume impact function for the developed agent-based model. The function is committed to computing price changes caused by the imbalance between the total demand and supply. The exact shape of the function may depend on various factors (Lillo *et al.* 2003). Among them: the selected time horizons where the impact is observed; the size of the traded volumes; the types of markets where the trades are performed and many others. Several research works have been done on this topic and different models were proposed. A stable and linear impact function was described by Kyle (1985). Huberman and Stanzl (2004) provided later the proof that permanent price impact must indeed be linear while the temporary one can be of a more general form. At the same time, more sophisticated non-linear functions were reported (Hasbrouck 1991, Hausman *et al.* 1992, Kempf and Korn 1999). We decided to choose the model proposed in the relatively recent work Bouchaud (2010) for our experiments. According to Bouchaud, the impact of trading volume is non-linear and one of the best approximations is the square root function. Therefore, we endow the market response to the agents aggregated trades in the following way:

$$r_n(\Delta N_n) \equiv S_n - S_{n-1} = \left\lfloor \alpha \mathrm{sgn}(\Delta N_n)\sqrt{|\Delta N_n|} \right\rfloor, \tag{2}$$

where: $r_n(\Delta N_n)$ is the price change (return) at the step n; α is the parameter limiting the minimum price shift; sgn(.) and $\lfloor . \rfloor$ are the sign and the floor functions correspondingly.

There is a situation among all possible scenarios when the number of agents deciding to flip short positions is less than the number of agents deciding to do the opposite just by one agent. This case leads to the imbalanced volume equal to the sum of the volume to close a position and the volume to open an opposite. Therefore, the smallest possible non-zero value ΔN_n is equal to two. The parameter α was chosen to guarantee that in such situation the market will respond to the smallest possible imbalance by changing the price by one unit: $\alpha = \sqrt{2}/2$. We use the floor function to simulate only discreet price changes typically observed in real markets.

Periods of economic equilibrium happen when demand meets supply (Debreu 1987). In this case, the sequence of trades happened at the market does not cause any price change, so the evolution of the price curve becomes temporarily stable. Such states do not last notably long. Even small market fluctuations enliven further asymmetric activity of the sell and buy sides. Eventually, the disturbance becomes portrayed in a sequence of substantial price moves. Similar to the reality, the proposed agent-based model is capable of producing zero difference between the number of all buyers and sellers which entails zero net volume. This volume being put into equation 2 does not cause any price change. The agents do not receive any new information and thus do not have a chance to make decisions on further trades without price changes. As a result, this puts the following evolution of the price curve on hold. To reactivate the trading, we add a small random price shift upward or downward with equal probability 0.5. The shift happens whenever the net volume ΔN_n is equal to zero. The size of the random move was chosen to be big enough to trigger a new intrinsic event for agents who have one or both thresholds equal to the size of the smallest price tick: $\delta = 1$. This way, these agents receive a new piece of information and try to flip their opened positions creating new demand or supply which leads to the consequent price change.

We used the next parameters for the main simulations: initial price level $S_0 = 0$; minimum price step (a tick) $r^{min} = 1$; $\alpha = \sqrt{2}/2$; smallest and biggest thresholds $\delta_{min} = 1$ and $\delta_{max} = 50$; the step between two consecutive thresholds on the grid equal to 1; total number of trading agents is 2500. Prices, produced by the model, are aggregated returns computed as responses to all imbalanced volumes that happened in the past. Therefore, the zero initial price simply indicates that no returns happened before the model was initialised. The smallest threshold $\delta_{min} = 1$ guarantees that any elementary price move will trigger a new intrinsic event for agents operating at the minimal scale. The size of the probability to flip position P_{flip} coincides with the empirically and theoretically found probability to register a new overshoot event before the directional-change one (Golub *et al.* 2014):

$$P_{flip} = e^{-1} = \mathbb{P}(\omega(\delta) \geq \delta). \tag{3}$$

The maximum possible price move r^{max} can be observed only when all agents defined in the model decide to flip their positions simultaneously. Thus, the largest absolute price change is determined by the number of agents on the grid and is directly connected to its length L:

$$r^{max} = \alpha L\sqrt{2}. \tag{4}$$

To generate such big return, all agents should register intrinsic events at the same moment of time and their opened positions have to be of the same type. In the real markets, such significant returns with negative sign are usually interpreted as crashes. They do not happen on a daily basis and are observed when large numbers of market actors accidentally make similar trading decisions in

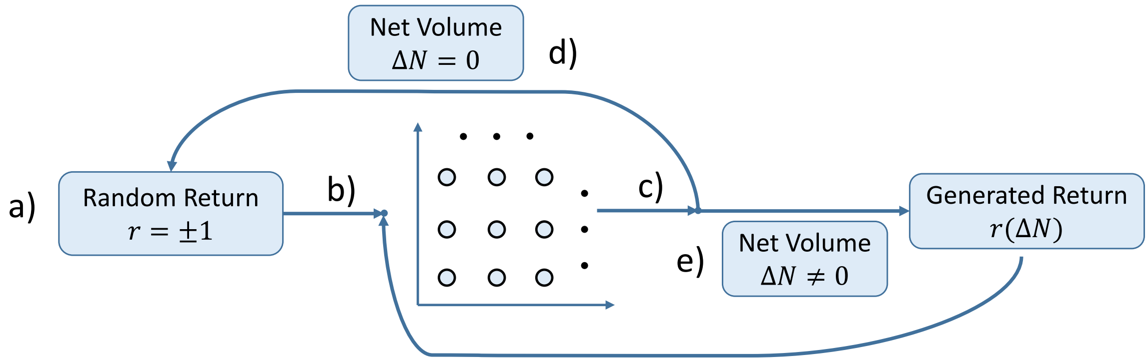

Figure 4. The schematic summary of the algorithm used to generate the synthetic time series. a) The model generates a random return in order to create the first signal for the group of agents represented by a grid of thresholds. b) All agents consider the latest price return as a signal in their local directional-change intrinsic time defined by the pairs of thresholds $(\delta_{up}, \delta_{down})$. Some group of the agents flip their opened positions reacting to the observed return. The generated net volume ΔN is expressed as the overall difference between the total number of agents decided to flip long and short positions. The net volume can be: d) equal to zero which means that no new signal is generated and a random price shift should happen to provide a new signal for the agents; e) not equal to zero which means the volume impact function (2) can generate a new return as the reaction to the aggregated traders behaviour. The new return is the input for the grid of agents at the next step of the simulation.

one critical instant. In our model, the probability to observe the maximum price move of the size r^{max} is likewise possible but minuscule. It is guaranteed by the condition that the decision to flip positions should happen at the same time for all agents. The latter happens with the probability $P_{flip}^{L^2}$ which rapidly tends to zero with increasing the size of the grid L.

The schematic summary of the algorithm used to generate the synthetic time series is provided in figure 4.

4. Benchmarks

The main goal of this research work is to check whether the agent-based model operating in the directional-change intrinsic time is capable of generating synthetic time series, statistical properties of which are coherent with the ones typical for foreign exchange market. Several benchmarks have been chosen to verify the accuracy of the model. The whole set of tests consists of four traditional methods often used in research with the same intent and a new approach wholly based on the directional-change intrinsic time. We propose that the latter is the best way of evaluating agent-based models. Further, we describe all selected benchmarks in details.

4.1. *Traditional benchmarks*

One of the well-known facts about the market microstructures is that price returns in any liquid market do not exhibit significant linear autocorrelation (Arneodo *et al.* 1996). This phenomenon is also formulated as the 'efficient market hypothesis': prices instantly and fully reflect all available information making it impossible to build a simple trading strategy based on the 'statistical arbitrage' (Basu 1977). Only at ultra-short time interval price curves could have slightly correlated historical returns. At this scale, the market, as a global multi-agent system, is still absorbing a new piece of information. According to Cont *et al.* (1997), 'in a few minutes' the autocorrelation can be safely assumed to be equal to zero. Low autocorrelation of returns is one of the most popular stylised facts which regularly accompany liquid markets. Therefore, we selected that statistical property to be among our benchmarks. The autocorrelation function of the series of returns M

with mean μ and variance $\sigma(\tau)$ computed using the given lag τ is defined as

$$R(\tau) = \frac{\mathbb{E}[(M_t - \mu)(M_{t+\tau} - \mu)]}{\sigma(\tau)^2}. \tag{5}$$

The second stylised fact chosen for the analysis is related to the fat-tailed distribution of returns at the frequency higher or equal to one day. This property, also known as excess kurtosis, is thoroughly discussed in the book Mandelbrot and Hudson (2010). The authors point out that there is the full range of theories build on top of the assumption that returns are normally distributed (including the famous work of Black and Scholes (1973)). Nevertheless, the processes in the real financial markets have always been different from Brownian motion and this assumption is a severe flaw of any related financial model. Therefore, it is important to create the model which is able to reproduce returns characterised by the fat-tailed distribution. We measure the excess kurtosis k in the following way:

$$k = \frac{\langle (r(t,\tau) - \langle r(t,\tau) \rangle)^4 \rangle}{\sigma(\tau)^4} - 3, \tag{6}$$

where $\sigma(\tau)$ is the variance of log-returns $r(t,\tau) = S(t) - S(t-\tau)$ computed with the lag τ. Excess kurtosis $k = 0$ means the absolutely normal distribution. Values bigger than zero indicate deviations from it. Brown and Warner (1985) demonstrated that in the stock market excess kurtosis is usually below 4; Cont (2001) mentioned that for S&P 500 futures the value is around 16; the kurtosis is equal to 60 for Dollar/Swiss Franc futures; in Gençay *et al.* (2001) one can find that for USD FX rates it is roughly 30 (through 10 minutes lag intervals).

Excess kurtosis is typically positive when time lag is relatively small and it tends to zero as the lag increases. This fact is called the aggregated normality or aggregational Gaussianity and can be accounted for by the 'mixture of normals' (Antypas *et al.* 2013). We select several increasing time lags to demonstrate that the distribution of returns produced by the agent-based model has distinct excess kurtosis. The values of the kurtosis are expected to be dependent on the selected scale. Thus, the aggregated normality is the third benchmark in our set.

The scale-invariance of the absolute price change (return) to the period of time when it occurs (Müller *et al.* 1990, Mantegna and Stanley 1995, Dacorogna *et al.* 2001) is the fourth selected benchmark. There is no agreement on the origin of this invariance called scaling law or power law and various assumptions can be found in the literature (Bouchaud 2001, Farmer *et al.* 2004, Joulin *et al.* 2008). Its omnipresence in finance incentivised scientists to apply it for real financial problems such as risk management and volatility modelling (Ghashghaie *et al.* 1996, Gabaix *et al.* 2003, Di Matteo 2007). We check whether the returns generated by the directional-change agent-based model are also characterised by this power law.

Most of the scaling laws found in financial markets rely on intervals measured in physical time. It is the essential fact to be aware of constricting any model of the real financial system. As it was discussed in Section 2, measures of the markets dynamic performed in physical time are not efficient in catching extreme events. Physical time does not adjust its flow to the speed of actions affecting the market. The extreme events, in turn, are the most critical information for statistical analysis. On the other hand, agent-based models are neither capable of 'feeling' the flow of physical time (they are a piece of code, after all). Therefore, it is not a trivial question of how one can connect the sequence of actions performed by simulated agents to real seconds, days and years. For example, the scale of volatility has a direct impact on the number of observed directional-change intrinsic events (see figure 1). Each intrinsic event is a signal for the corresponding agent who registered it. Thus, the volatility clustering means compression and expansion of the inner clock used by the agent-based model. The exact way of how the agents activity relates to the speed of physical time should be algorithmically predetermined. The assumption that the agents make decisions at equidistant moments of time, for example, every second, is the most popular approach used to bridge the gap between physical and intrinsic time in agent-based modelling. In this case, for instance, 20 000 000 steps in intrinsic time would correspond to 231.5 trading days (that is, close to a full trading year or 252 days). To validate the model by the traditional benchmarks, we follow the same principles and postulate that price changes can be observed only over discrete moments of time.

4.2. *The ultimate benchmark*

Scaling laws are ubiquitous properties of our world and are present in any domain of natural and social phenomena such as physics, biology and social sciences (Andriani and McKelvey 2007). The goal of the agent based model is the simulation of a collection of returns statistical properties of which coincide with the ones typical for the real financial markets. Therefore, it is also important to consider such omnipresent scaling properties while validating artificial sets of interacting agents.

The traditional benchmarks mentioned in the previous section can be successfully applied for describing processes and actions happening in the real world. Time is the universal measure of actions and interactions in such real environment. Therefore, all these benchmarks account for physical time as the factor of scale. Unfortunately, the artificial agents have no 'feeling' of elapsed periods and operate using the 'signal-reply' logic only. Therefore, applying these benchmarks as well as methods and tools where time plays the crucial role for interpreting the agents automated actions is not fully correct. We consulted the work of Glattfelder *et al.* (2011) in the search for a new omnipresent property of high-frequency markets which can be ultimately used to verify any agent-based model. The main criteria: it should be possible to employ the property as a benchmark independently on the relation of physical time and the sequence of agents actions. Fortunately, there is one, out of 12, newly described scaling laws which perfectly fits the mentioned criteria: the 'overshoot scaling law'. It is fully based on the directional-change concept where only relative price moves dictate sequences of events dismaying time intervals between them. According to Glattfelder *et al.* (2011) the average length of the overshoot section is approximately equal to the size of the directional-change threshold. Golub *et al.* (2014) analytically showed that in the continuous process with zero trend the probability of overshoot $\omega(\delta)$ reaching the length l equals to $\exp(-l/\delta)$, i.e.

$$\mathbb{P}(\omega(\delta) = l) = \exp\left(-\frac{l}{\delta}\right). \tag{7}$$

This dependence reveals the exponential relation between the length of a directional-change threshold δ and the expected length of the overshoot $\omega(\delta)$. From equation (7) it follows that the expected overshoot $\mathbb{E}[\omega(\delta)]$ is equal to the size of the directional change threshold δ:

$$\mathbb{E}[\omega(\delta)] = \delta. \tag{8}$$

Glattfelder *et al.* (2011) have empirically shown that the average coefficients E_{OS} and C_{OS} of the overshoot scaling law $\langle\omega(\delta)\rangle = \left(\frac{\delta}{C_{OS}}\right)^{E_{OS}}$ across all 13 currency pairs analysed in the work are $E_{OS} = 1.04$ and $C_{OS} = 1.06$. Lowest and the highest registered values specific for the pairs are: $C_{OS,low} = 0.98$, $C_{OS,high} = 1.17$; $E_{OS,low} = 0.98$, $E_{OS,high} = 1.08$. This statistical property is fully agnostic to the volatility and does not relate on any values defined in physical time. Its application as a benchmark to an agent-based model does not require any additional assumption on the connection of agents activity and the flow of real physical time. Therefore, we call that scaling law as the main benchmark of our model.

We used the same notation proposed in Glattfelder *et al.* (2011) to validate scaling laws mentioned above. We build log-log plots where on X- and Y-axis the base and dependent values of scaling laws are placed. We assume a linear relationship between the response variable Y (for example, the average size of price returns) and the fixed variables X (for example, a period of time): $Y = A+BX$, where A and B are unknown parameters to be estimated. Thus, B defines the slope of the line on

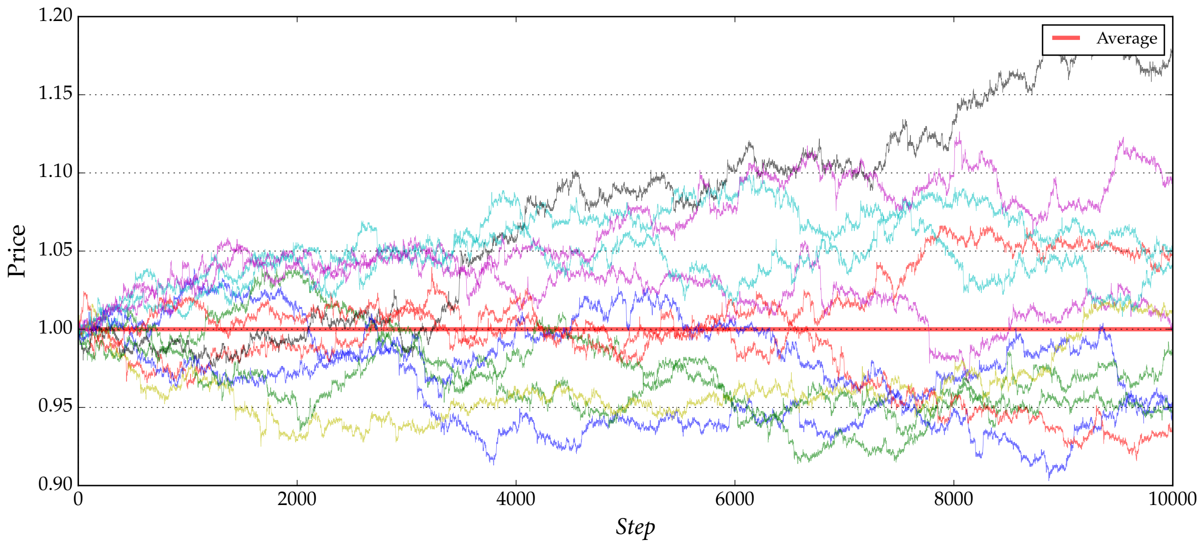

Figure 5. 10 time series generated by the agents from the entire grid of 50 by 50 equally distributed points (see figure 3). The presented price curves were obtained by applying the exponential function to the set of logarithmic returns generated by the model: $Price = \exp(S_{log}/1000)$. Here 1000 is just a scaling factor.

the log-log plot and A is the intersection of Y-axis. In this way, a scaling law takes the following form:

$$y = \left(\frac{x}{C}\right)^{E}, \tag{9}$$

where $y = \exp Y$, $x = \exp X$, $E = B$ and $C = \exp(-A/B)$.

5. Results

In this chapter, we highlight the main findings of the research work and review how components of the agent-based model affect the properties of the generated time series.

All experiments have been performed in two steps. First, we analysed time series generated by the agents from the entire grid operating simultaneously; second, we examined the effects induced by I or III part of the grid being performing solely. The following two sub-chapters outline details of each experiment.

5.1. *The entire grid*

As the first step, we tested the performance of the whole grid of directional-change agents. Parameters of the experiment are specified in Section 3. An example of 10 price curves generated by the model with the help of the squared root impact function (2) is presented in figure 5. The curves consist of various intervals with plateaus and sudden jumps thus mimicking features of the real FX market. At the same time, there is no any prevailing trend which would define the evolution of each given price curve. The red line represents the average aggregated return computed as the result of 1000 independent simulations. This line remains horizontal throughout all 10 000 steps. The latter confirms that the model where all agents participate in the trading does not possess any deterministic impact on the trend.

Benchmarks introduced in Section 4 were applied to validate the synthetically generated data sets. Autocorrelation function (ACF) of a time series consisting of 10 000 000 simulated returns is

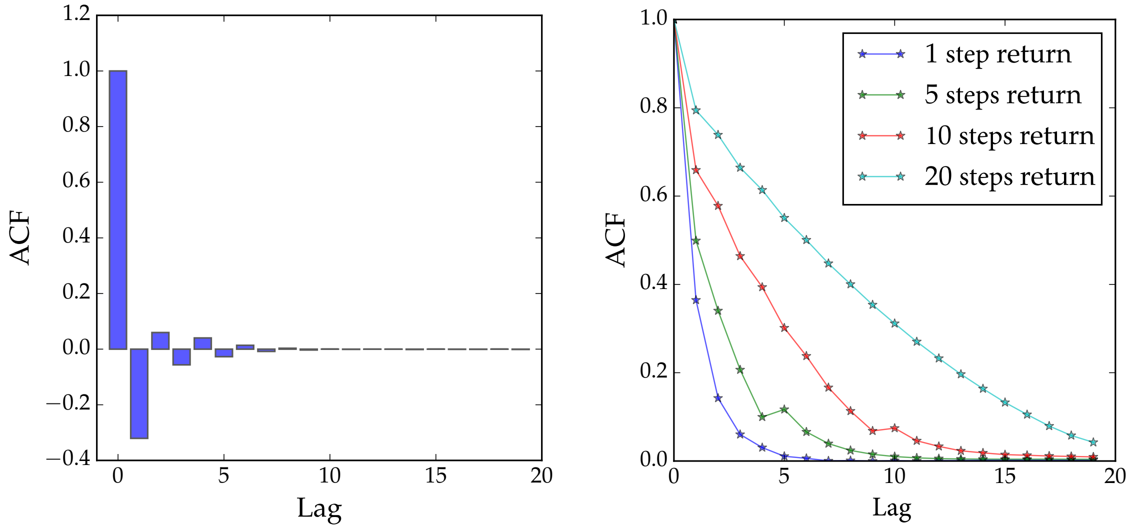

(a) (b)

Figure 6. (a) Autocorrelation function of the generated time series. Lags are measured in the numbers of elementary price moves. (b) Autocorrelation function of absolute returns computed on the same time series using four different lag steps.

shown in figure 6(a). The maximum negative correlation (-0.32) is observed for the lag size equal to one step. The rest of the autocorrelation values is significantly smaller. The autocorrelation function rapidly decays and becomes indistinguishable from zero after about ten steps.

Figure 6(b) shows slowly decaying autocorrelation functions of absolute returns computed over four different lag steps. The bigger the lag step, the lower the decline of the ACF. This effect confirms the ability of the agent-based model of replicating the slow decay of autocorrelation in absolute returns. That is, the aggregational normality has been successfully validated.

The absolute price change scaling law is shown in figure 7(a). C and E are the scaling parameters from equation (9) described in Section 4.2. R is the Pearson product-moment correlation coefficient. One can clearly observe the linear dependence of the absolute price change and the elapsed time interval expressed in logarithmic values. It is important to highlight that this scaling law employs physical time to measure the intervals between given prices of the time series. There is no academic agreement on the periodicity and intensity of trades assumed to be accomplished by artificial agents in physical time. Both these values are determined by the assumption made in every particular model. High trades frequency induced big size of price changes over short periods of time. Low frequency results in the opposite. That is equivalent to the compression of time for the same agents activity level. It leads to the increased or decreased volatility of the process and the bigger or lower slope of the log-log plot $Y = A + BX$ (see Section 4.2 for more details). The scaling parameters E and C are directly connected to the coefficients of the plot through equation (9). Therefore, the choice of the trade frequency in physical time directly affects the obtained scaling coefficients of the absolute price change scaling law. As results, it is not possible to compare the parameters observed in the FX markets with the ones generated by the model. Nevertheless, the fact that the absolute price change of generated returns depicted in figure 7(a) is represented by a straight line on the log-log plot (described by a power law function) is enough to confirm the successful replication of the absolute price change scaling law.

Figure 8 shows distributions of returns generated by the model. The returns are computed over four gradually increasing lags: 10, 50, 100, and 1000 steps. As it can be seen from the figure, the distributions with lags up to the hundreds of steps (figures 8(a), 8(b), and 8(c)) are characterised by noticeable fat tails. The fat tails disappear around the lag level of 1000 steps (figure 8(d)). Assuming that the agents generate a new return every minute, the observed effect is in line with

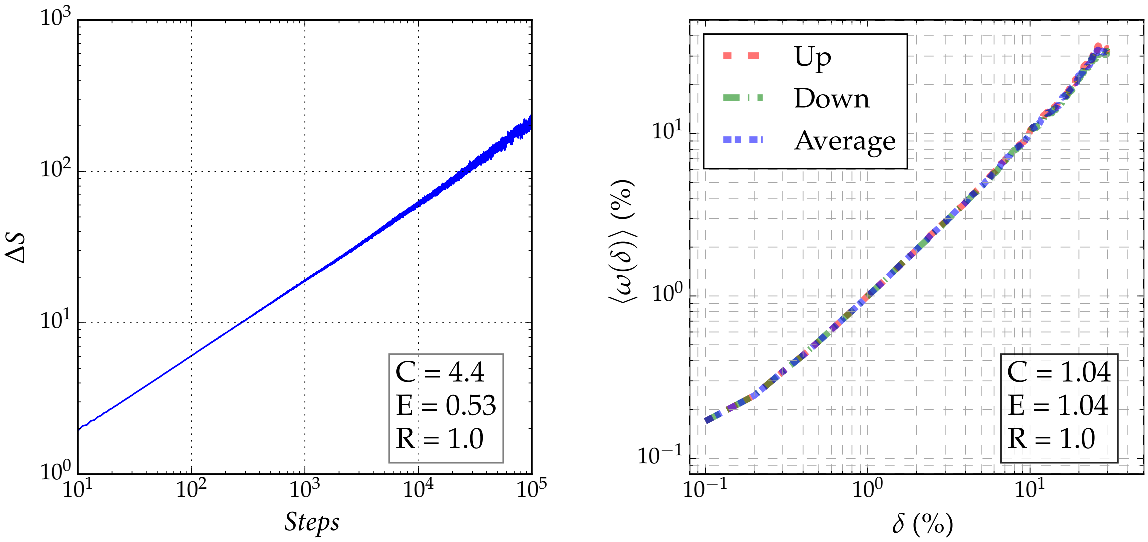

(a) (b)

Figure 7. (a) Average absolute price move as the function of the number of steps (the absolute price move scaling law). (b) Overshoot scaling law build from a time series generated by the agents from the entire grid. Parameters on the plot correspond to the average line. Approximation was done for $\delta > 0.3\%$. The same equation provided in the description of figure 5 was used to transform thresholds from absolute values to relative ones (into percentage terms). Coefficients of the Up line: $C = 1.04, E = 1.05, R = 1.0$; of the Down: $C = 1.03, E = 1.03, R = 1.0$. The plot is based on 20 000 000 steps.

empirical results found in real markets (for example, values from Kullmann *et al.* (1999)). Measured excess kurtosis k is equal to 2.73 at 10-steps lag and only 0.06 when the lag rises to 1000 steps. The excess kurtosis rapidly decreases together with the growing lag size confirming the empirically observed dependence. We present two probability plots for lags of 10 and 1000 steps in figure 9. The probability plots confirm the aggregated normality once again. Nevertheless, the empirical analysis of Kullmann *et al.* (1999) shows that the excess kurtosis at such small lags is usually much higher ($k = 10$ and more). The intrinsic time agent-based model can also be used to generate returns characterised by the size of the kurtosis typical for the real price data. The size of the grid and thus the quantity and the diversity of trading agents directly contribute to the excess kurtosis. The grid 50 by 50 points produces returns with $k = 2.73$ at 10-steps lag; the grid 100 by 100 points result in $k = 3.46$; the grid 200 by 200 points generates time series with $k = 6.01$.

Finally, we validated the agent-based model using the properties of the overshoot scaling law. The log-log plot of the average overshoot length versus the size of the directional-change threshold was constructed (figure 7(b)). The revealed dependence appears to be linear. The scaling parameters are $E = 1.04$ and $C = 1.04$. These values are exceedingly close to the ones observed in the real foreign exchange market (see Section 4.2). We also investigated two supplementary versions of the overshoot scaling law in addition to the one described in Glattfelder *et al.* (2011): the dependence computed using only overshoots following upward directional changes (red dashed line in the figure) and only overshoots following downward directional changes (green dashed line). There is no noticeable difference between all three lines. This observation confirms that the directional-change agent-based model is able to reproduce the overshoot scaling law. Results of the experiment are also in line with the theoretical equations (A6) and (A6) provided in Appendix A.

Markets react to exogenous shocks such as news in multiple ways (Frank and Sanati 2018). Positive and negative news could as permanently change the preceding price trend as well as make a short-term disturbance. The latter can become rapidly absorbed by the system or become pronounced in the corresponding price change. This sort of price change might be also observed as instantly after the respective event as well as over significantly longer period. We checked how

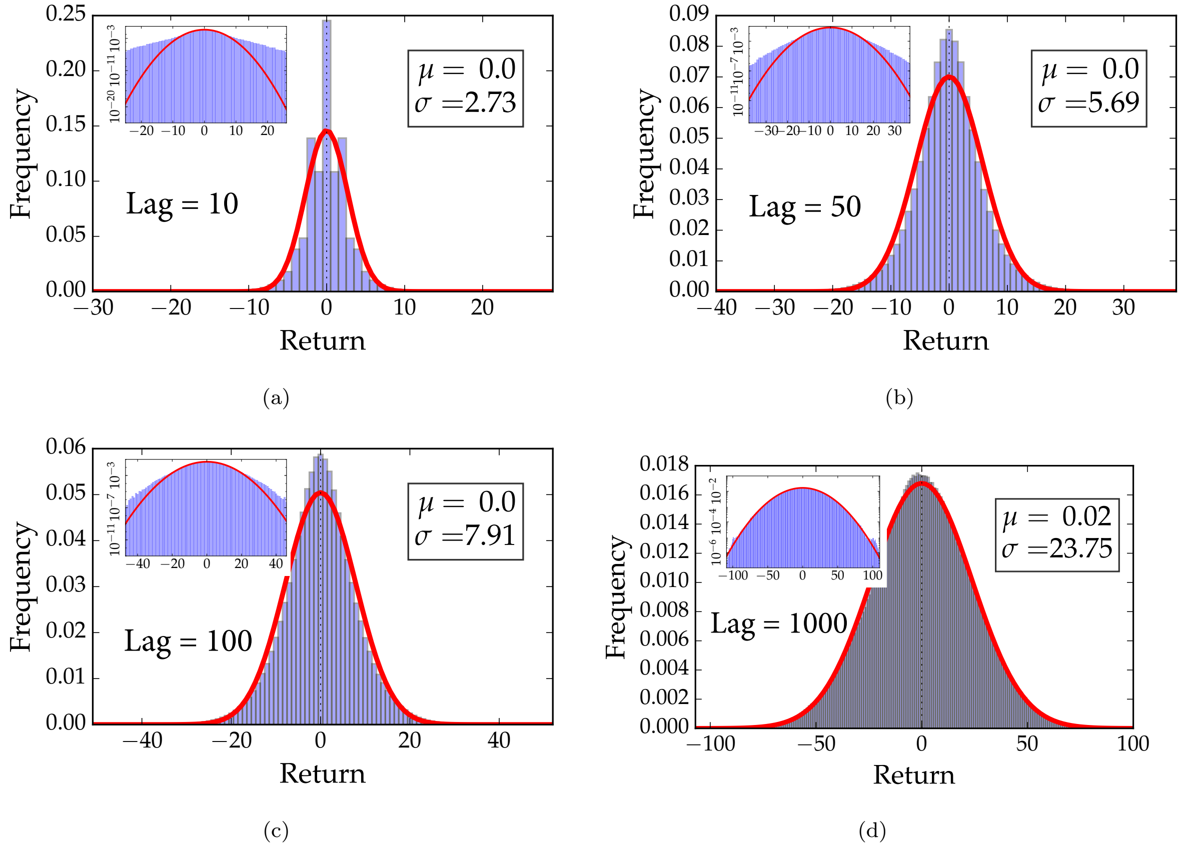

Figure 8. Distribution of returns generated by the agents from the entire grid and computed at four different lags. Their approximation by the Gaussian function is marked by red lines. Included sub-plots are the same distributions but presented in the logarithmic scale.

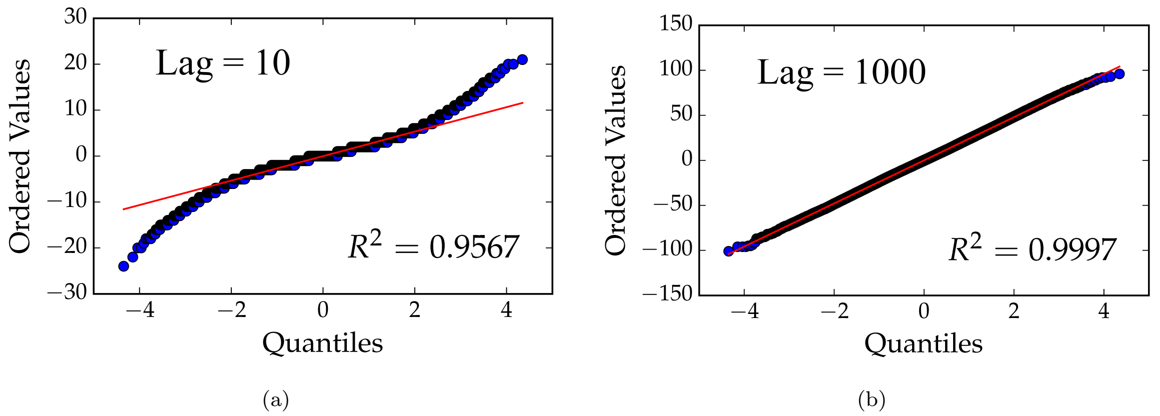

Figure 9. (a) Probability plots of returns computed with the lag equal to 10 steps. (b) Probability plots of returns computed with the lag equal to 1000 steps. The total number of prices in the generated time series is 10 000 000.

the agent-based model operating in intrinsic time reacts to instant shocks of various amplitudes. Each shock was simulated by a small portion of extra volume on the sell or the buy side added to the net volume. The Monte Carlo simulation demonstrated that the extra volume added to the net volume at the arbitrary step has instant and permanent impact on the average return. The size of the impact is defined by equation (2). Results of the simulation can be requested from the authors.

Volume impact function connects the net volume generated by the set of agents at time t with the induced price change observed at $t+1$. The square root dependence (2) was selected for the agent-based model although there is no scientific agreement on the universal form of the functions. Therefore, linear and logarithmic impacts have also been tested in our work. Time series generated with the help of these dependencies do not excel the same quality of stylised facts reproduced by the square root function. The linear impact induces rapid price fluctuations around the initial level and none of the statistical properties become observed. In turn, time series generated using the logarithmic dependence expectantly have similar properties to the ones of time series produced with the help of the squared root. The fat-tailed distribution of returns is of the same scale for both functions. Nevertheless, scaling coefficients of the overshoot scaling law become noticeably worse in case of the logarithmic impact. The coefficients are: $C = 0.86$, $E = 0.95$, $R = 0.16$ for the linear impact; $C = 1.38$, $E = 0.99$, $R = 1.0$ for the logarithmic impact; $C = 1.04$, $E = 1.04$, $R = 1.0$ for the squared root function (2).

The linear distribution of thresholds and the squared shape of the grid are aimed to replicate the diversity of real groups of traders in FX markets. The selected parameters only approximately replicate the composition of the trading strategies existent in the real world. Nevertheless, we experimented with different sets of thresholds, such as their logarithmic distribution as well as the radial fit and the circular shape of the grid. The obtained performance of the agent-based model was close to the one offered by the simple squared grid whenever the number of agents remains substantial.

5.2. *Asymmetric regions*

The entire grid of directional-change agents generate time series with zero average trend (Section 5.1). We claim that the reason of the observed phenomenon is the selected symmetry of the grid: to each agent from the region I with not equal thresholds δ^{I}_{up} and δ^{I}_{down} there is an agent in the region III with thresholds δ^{III}_{up} and δ^{III}_{down} such that $\delta^{I}_{up} = \delta^{III}_{down}$ and $\delta^{I}_{down} = \delta^{III}_{up}$. To confirm the statement, we performed a series of tests where the imbalance of trading agents was created. The imbalance is achieved by removing a part of agents from the initial grid. The effect observed in the test is the clearly pronounced non-zero average trend of the generated time series. We demonstrate results of two trivial experiments where only agents from the region I or region III have been selected to trade (figure 10). The induced deviations of the average trend are upward and downward correspondingly. The permanent trend observed in both experiments is present due to the created imbalance between the number of agents supporting trends of opposite directions. Agents from the region I tend to trade more often when the price is going down while agents from the region III prefer the rising trend.

Several factors affect the average slope of all generated price curves: the total number of agents forming the initial grid, the section of the grid used to generate a time series, the selected time interval between two consecutive steps. We performed a set of experiments to visualise the impact of the created grid asymmetry on generated prices. An extra directional-change agent with the specific pair of thresholds was added to the initial set of agents forming the 50 by 50 points grid. This resulted in $2500 + 1$ agents trading simultaneously. We show in figure C3 the distribution of final aggregated returns after 10 000 steps averaged over 1000 independent simulations. Additional agents with symmetric thresholds $\delta_{up} = \delta_{down}$ have no pronounced impact on the average aggregated return (parameter μ of the Gaussian distribution used for approximations) as well as on the standard deviation (parameter σ). Agents with asymmetric thresholds tend to deviate the average trend in the direction which coincides with the direction of their bigger thresholds. The deviation is bigger in case of smaller δ_{up} and δ_{down}. Agents with such thresholds are similar to high-frequency traders and eager to react to all trend changes. As result, they make more trades than the traders with substantially bigger directional-change thresholds (see figure 1(a)). These facts explain the observed phenomenon. A set of overshoot scaling laws computed using the data generated in each experimental setup is presented in figure C4. It is worth noting that lines related to the upward

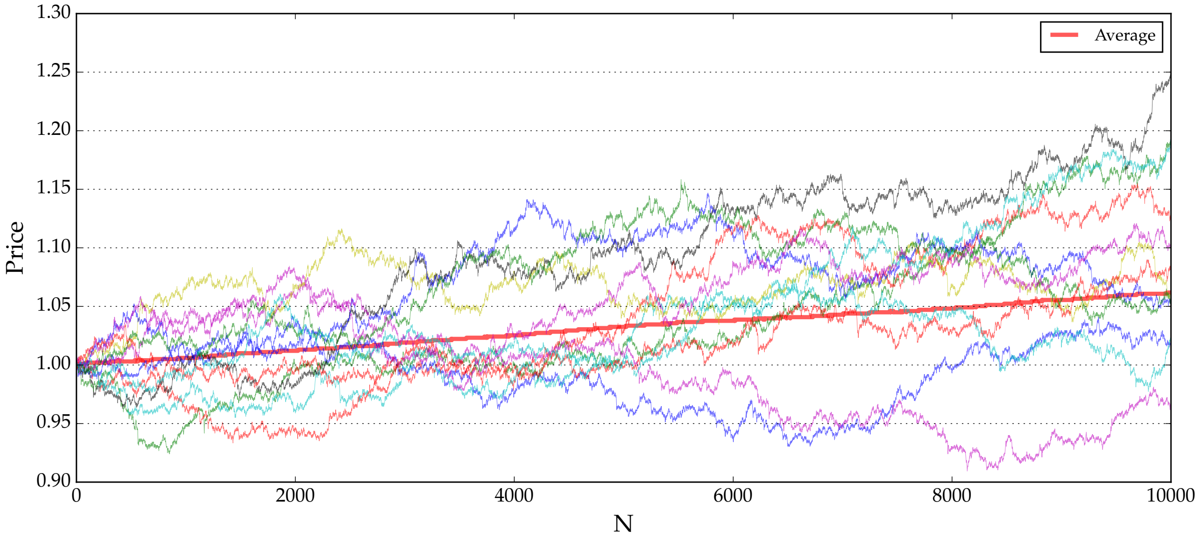

(a)

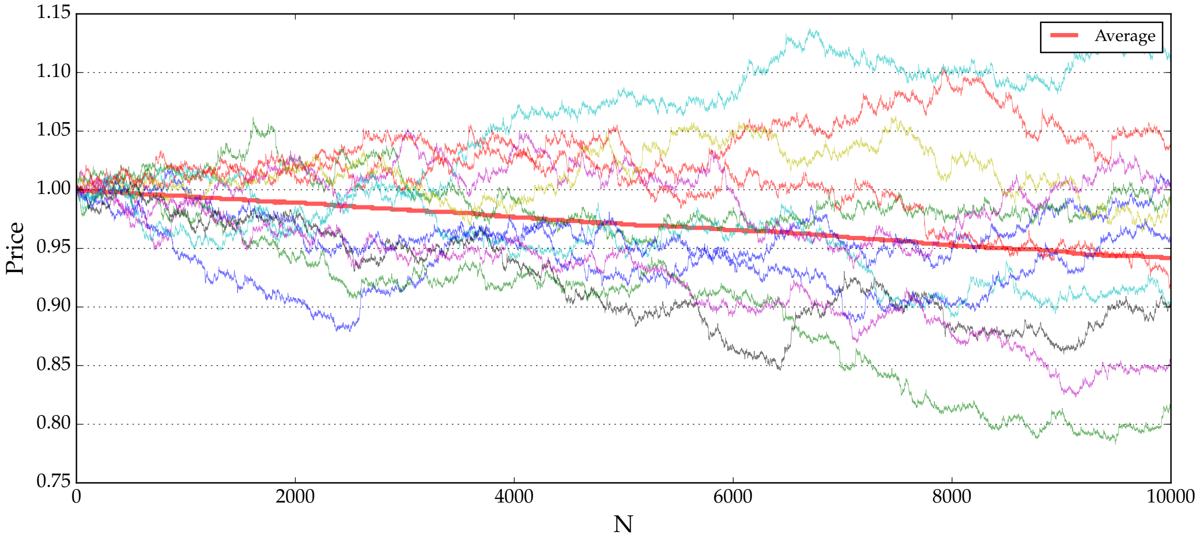

(b)

Figure 10. (a) Time series generated by agents from the region I of the initial grid and (b) by agents from the region III. Red lines represent average prices based on 1000 independent simulations. The initial grid size is 50 by 50 points.

and downward overshoots are symmetrically shifted from the diagonal line.

The number of agents participating in the trading effects the excess kurtosis of generated returns (see Section 5.1). Grids of different sizes were tested to investigate whether there is any other impact. Only agents from the I region were involved in the experiment. The average aggregated returns of 1000 independent simulations are shown in figure 1(b). As it can be seen from the figure, the bigger the size of the grid, the higher the deviation of the aggregated trend from the trend-less case. The net volume which is the only parameter of the impact function (2) is the reason of the observed deviation. Bigger number of agents creates higher standard deviation of the net volume. The latter induces the revealed in figure 1(b) dependence of the grid size and the average trend. The dependence has been also confirmed in another experiment results of which are presented in figure C5. The average positive and negative volumes generated by the model were computed as the function of the grid size. The experiment proves the statement that the bigger the size of the grid, the higher the generated absolute net volume.

We would like to note that the stylised facts selected for this project are still reproduced even when a part of the agents from the grid is deactivated. In figure C2 we include overshoot scaling laws computed using the data generated by two separate parts of the grid: by region I (figure 2(a)) and region III (figure 2(b)). The scaling coefficients C and E of studied dependencies insignificantly differ from the ones observed empirically in Glattfelder *et al.* (2011). Therefore, we claim that the agent-based model operating in directional-change intrinsic time can be used to simulate financial time series with statistical properties closely related to high-frequency markets. Moreover, the trend of the collection of generated returns can be specifically predefined. The set of the directional-change agents is the parameter which can be used to setup the trend. A comprehensive analysis of its precise direction is a topic for an independent work which is beyond the scope of this research paper.

6. Conclusion

An agent-based model where trades happen in directional-change intrinsic time was tested in this work. A set of artificial agents mimics the behaviour of real market participants by buying/ selling one lot of the traded asset. All trades happen at the moments when the directional-change intrinsic time ticks. Two types of intrinsic events are considered: directional changes and overshoots. Directional changes are moments when the price curve makes reversals equal to the size of the selected threshold. Overshoot intrinsic events happen every time when the overshoot length is a multiple of the of the corresponding directional change threshold. The agents flip their opened positions from long to short and vice versa at times when they observe intrinsic events of given magnitude. They flip the positions with the empirically and theoretically found probability to register a new overshoot event before the directional-change one. The probability is aimed to mimic behaviour of real market participants who do not exploit every trading opportunity. Traders decisions are not conditional on time intervals between registered intrinsic events. The agents are ignorant to the flow of physical time and consider prices as the only source of information about the markets activity needed to perform.

The agent-based model replicates a wide range of trading patterns observed in the real world. A unique pair of directional-change thresholds was assigned to each simulated trader. The size of the thresholds as well as their ratio define the scale of price trends analysed by each agent. The difference between the number of all agents decided to flip a long or short position at each new price represented the net volume or imbalance between the supply and demand. The net volume was used to calculate the impact of the imbalanced decisions on the further price change. The empirically observed squared root function connects the net volume and the corresponding market return.

A set of benchmarks was chosen to test the performance of the constructed agent-based model. The model is assumed to be successful if it manages to reproduce 'stylised facts' discovered in the real financial world. The utilised benchmarks: low auto-correlation of returns, fat-tailed distribution of returns, aggregated normality, price jump and the overshoot scaling laws. The last benchmark is fully agnostic to the connection of the simulated agents actions and the flow of physical time. We propose that the scaling law is the universal test for any agent-based model.

The presented agent-based model operating in the directional-change intrinsic time has successfully passed all chosen tests. The latter lets us make an educated guess that real market participants intentionally or unintentionally make trades using their own intrinsic time. In other words, they have preferences on the scale of price returns after which the traders reverse opened positions. The obtained knowledge can be used to substantially improve the quality of the inferences we make about the connection of the aggregated traders behaviours and observed price changes.

Appendix A: Overshoot as the function of trend

The average length of an overshoot section is approximately equal to the length of the corresponding directional change threshold (Glattfelder *et al.* 2011):

$$\langle \omega(\delta) \rangle \approx \delta. \tag{A1}$$

This dependence was found not only in historical tick data but also in arithmetic Brownian motion. Nevertheless, the average length of overshoots does not resemble the length of the selected threshold if the time series has a constant trend. In turn, the average overshoot length depends on the size of the trend (Golub *et al.* 2017). To find the exact form of this dependence, we select an arithmetic Brownian motion with increments dS_t, trend μ, and volatility σ:

$$dS_t = S_t - S_{t-1} = \mu dt + \sigma dW_t. \tag{A2}$$

Golub *et al.* (2014) derived the probability of the upward and downward overshoot to reach some fixed value x. The probabilities are the following:

$$\mathbb{P}(\omega(\delta_{up}) \geq x) = \exp\left\{ -\frac{x}{\sigma^2} \cdot \frac{(|\mu| - \mu) + (|\mu| + \mu)\exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}}{1 - \exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}} \right\}, \tag{A3}$$

$$\mathbb{P}(\omega(\delta_{down}) \geq x) = \exp\left\{ -\frac{x}{\sigma^2} \cdot \frac{(|\mu| + \mu) + (|\mu| - \mu)\exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}}{1 - \exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}} \right\}, \tag{A4}$$

where $\omega(\delta_{up})$ and $\omega(\delta_{down})$ stand for upward and downward overshoots correspondingly.

The expected value of the shown probability equations (A3) and (A4) $\overline{\mathbb{F}}(x) = \mathbb{P}(X \geq x)$ is

$$\mathbb{E}[X] = \int_0^\infty \overline{\mathbb{F}}(x)dx. \tag{A5}$$

Using (A5) one can find that:

$$\mathbb{E}[\omega(\delta_{up}, \mu, \sigma)] = \frac{\sigma^2 \left(1 - \exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}\right)}{(|\mu| - \mu) + (|\mu| + \mu)\exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}}, \tag{A6}$$

$$\mathbb{E}[\omega(\delta_{down}, \mu, \sigma)] = \frac{\sigma^2 \left(1 - \exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}\right)}{(|\mu| + \mu) + (|\mu| - \mu)\exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}}. \tag{A7}$$

The expected length of the all overshoots is the average of upward and downward expected overshoots (equations (A6) and (A7)):

$$\mathbb{E}[\omega(\delta_{up}, \delta_{down}, \mu, \sigma)] = \frac{\sigma^2}{2} \left(\frac{1 - \exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}}{(|\mu| - \mu) + (|\mu| + \mu)\exp\left\{-\frac{2|\mu|\delta_{up}}{\sigma^2}\right\}} + \frac{1 - \exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}}{(|\mu| + \mu) + (|\mu| - \mu)\exp\left\{-\frac{2|\mu|\delta_{down}}{\sigma^2}\right\}} \right). \tag{A8}$$

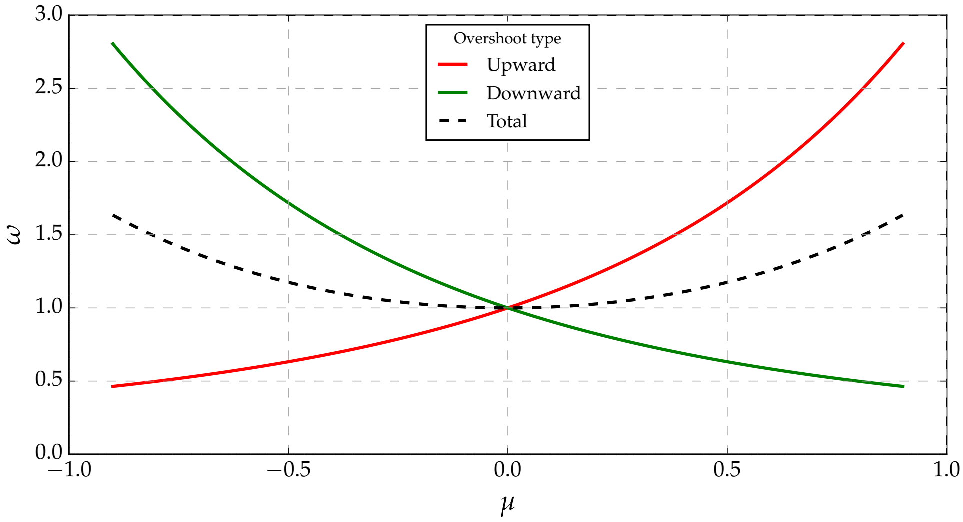

Figure A1. Expected length of the overshoot section ω as the function of trend μ computed using equations (A6) and (A7). The variance σ and threshold δ are fixed and equal to 1.

In figure A1 we demonstrate the dependence of the overshoot length ω on various trends μ when volatility σ and the threshold size δ are fixed. The last two variables are assumed to be equal to one ($\sigma = 1$; $\delta = 1$). It is easy to notice that the lengths of upward and downward overshoots coincide with each other only in case of zero trend. One can observe significant divergence of the curves for any other value not equal to zero. The obtained result is quite intuitive: price tends to overshoot more after an upward directional-change event in case of the overall ascending trend. At the same time, it tends to make a directional-change reversal right after a downward directional-change event. This observation suggests that for each non-zero trend μ there are such thresholds δ_{up} and δ_{down} that the total number of overshoot intrinsic events registered in the given time series is equal to the number of overshoots observed within a time series of the same length and zero trend. This property was directly used in the proposed agent-based model to replicate the diversity of trading patterns (Section 3).

Appendix B: Dissection Algorithm

By S_{tick} we mark the latest observed price, by S_{ext} the local extreme, mode is the current mode of the alternating trend which can be equal either *up* or *down*, δ_{up} and δ_{down} are upward and downward thresholds respectively, S_{IE} is the price at which the latest intrinsic event was observed. The algorithm returns 1 and -1 when the price curve hits levels of upward and downward directional-change events correspondingly. 2 and -2 will be returned in case of overshoot intrinsic events registered on ascending or descending trends.

Algorithm 1 Intrinsic Event

if *first tick* **then**
  $S_{ext} \leftarrow S_{tick}$
  $S_{IE} \leftarrow S_{tick}$
  return 0
else if *mode* is *up* **then**
  if $S_{tick} - S_{ext} \geq \delta_{up}$ **then**
    $mode \leftarrow down$
    $S_{ext} \leftarrow S_{tick}$
    $S_{IE} \leftarrow S_{tick}$
    return 1
  else if $S_{tick} < S_{ext}$ **then**
    $S_{ext} \leftarrow S_{tick}$
    if $S_{IE} - S_{ext} \geq \delta_{down}$ **then**
      $S_{IE} \leftarrow S_{tick}$
      return -2
    else
      return 0
else if *mode* is *down* **then**
  if $S_{ext} - S_{tick} \geq \delta_{down}$ **then**
    $mode \leftarrow up$
    $S_{ext} \leftarrow S_{tick}$
    $S_{IE} \leftarrow S_{tick}$
    return -1
  else if $S_{tick} > S_{ext}$ **then**
    $S_{ext} \leftarrow S_{tick}$
    if $S_{ext} - S_{IE} \geq \delta_{up}$ **then**
      $S_{IE} \leftarrow S_{tick}$
      return 2
    else
      return 0

Appendix C: Additional experiments

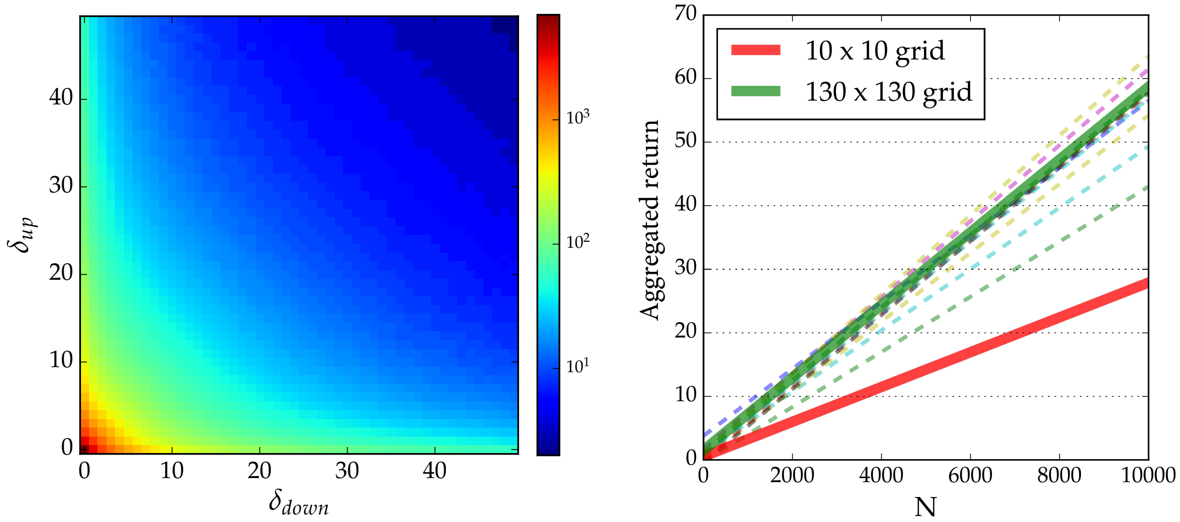

(a)

(b)

Figure C1. (a) The average number of trades performed by directional-change agents with parameters δ_{up} and δ_{down} within 10 000-steps simulation. Parameters of the agents used for the simulation are provided in Section 3. (b) Average aggregated returns recorded within 10 000 steps (N) of 1000 independent simulation when only agents from the region I are active. 13 tests with grids of various sizes were performed. The smallest grid is 10 by 10 points. The biggest is 130 by 130. The applied increment is 10.

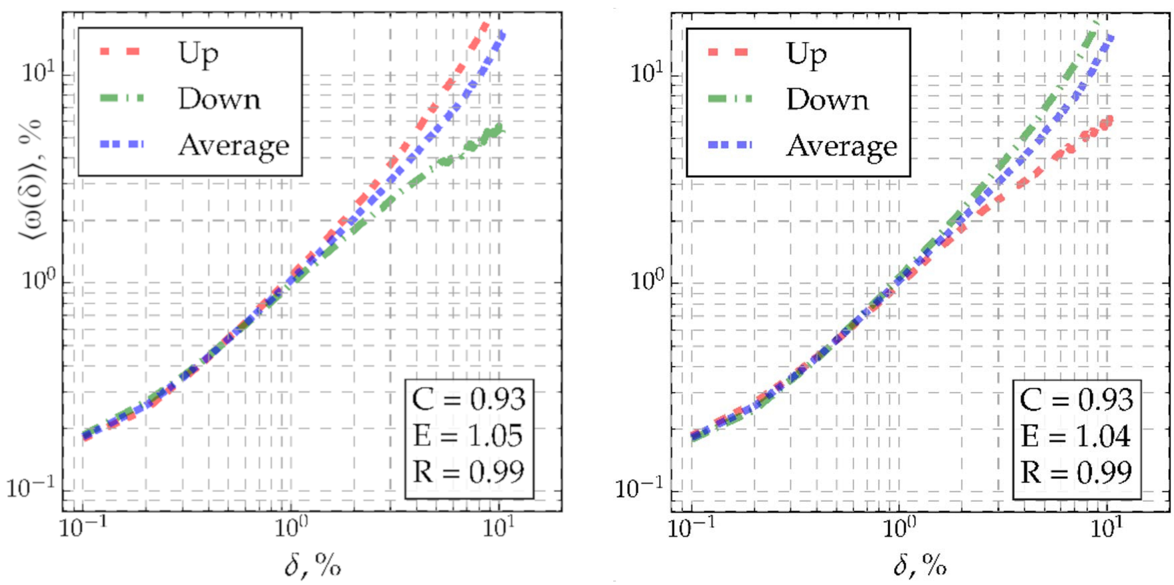

(a)

(b)

Figure C2. Overshoot scaling laws computed using time series generated by the agents from (a) I and (b) III regions of the grid only. Red dashed line (Up) corresponds to the average overshoot computed after directional-changes upward. Green dashed line (Down) corresponds to the average overshoot measured after directional-changes downward. The scaling parameters C, E, and R are shown for the lines representing the the average overshoot lengths (blue dashed line).

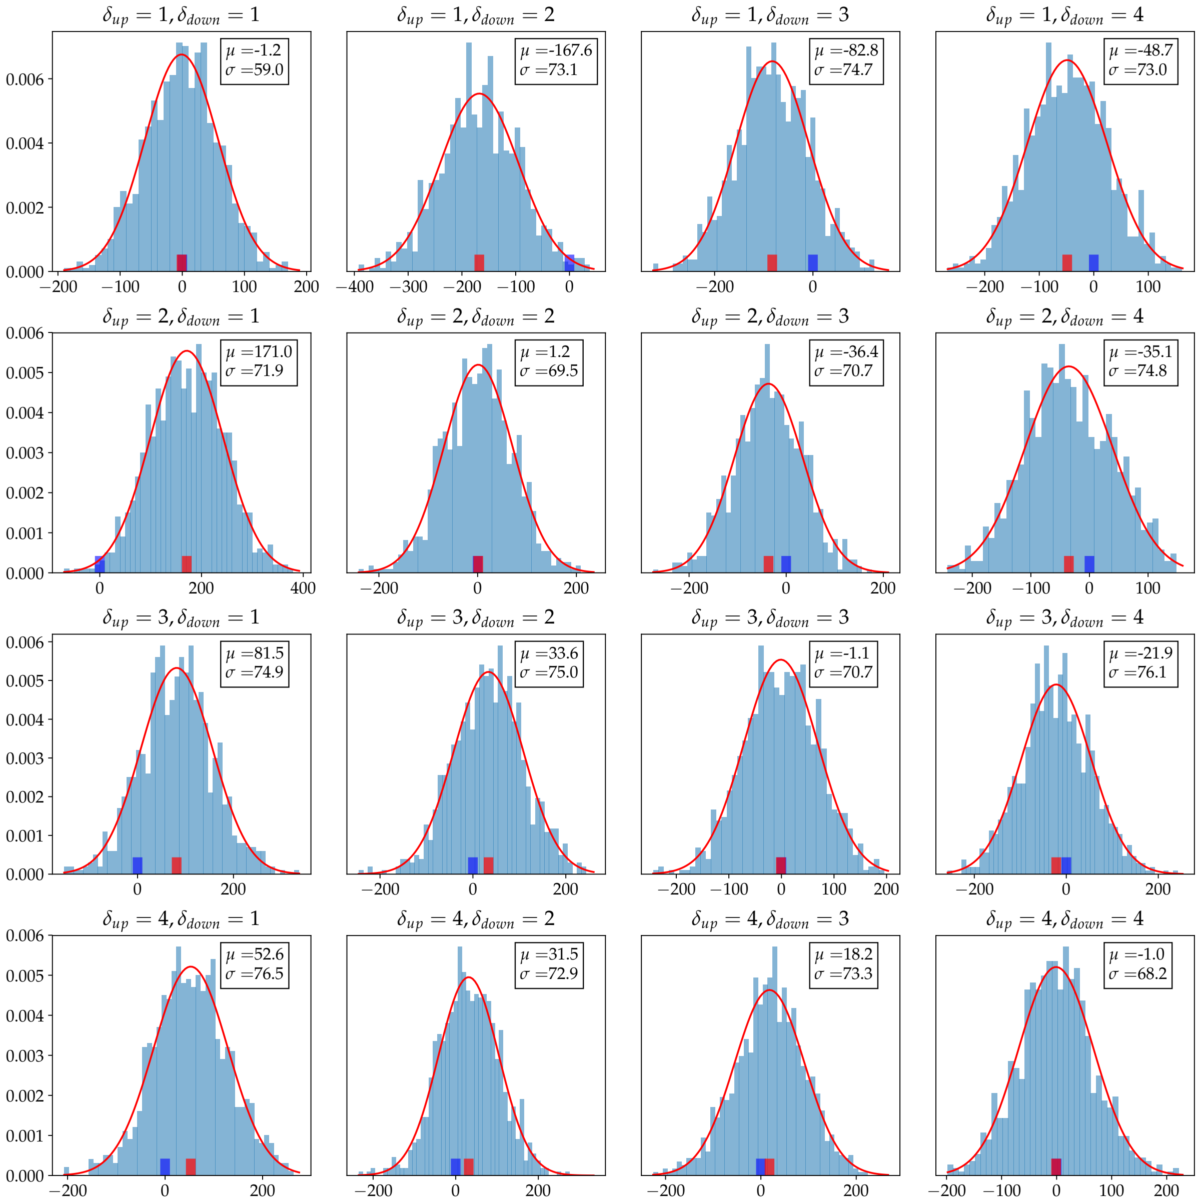

Figure C3. Impact of an additional intrinsic time agent on the average trend generated by the entire grid. $2500 + 1$ traders participate in each simulation. Parameters of the supplementary agents are put on top of each subplot. Experiment consists of 1000 independent simulations and 10000 steps. Normal distributions (red lines) approximate the final aggregated returns. Blue rectangular mark zero trend level. Red rectangular stand for the centre of the obtained distributions. The subplots are centred on the average aggregated return.

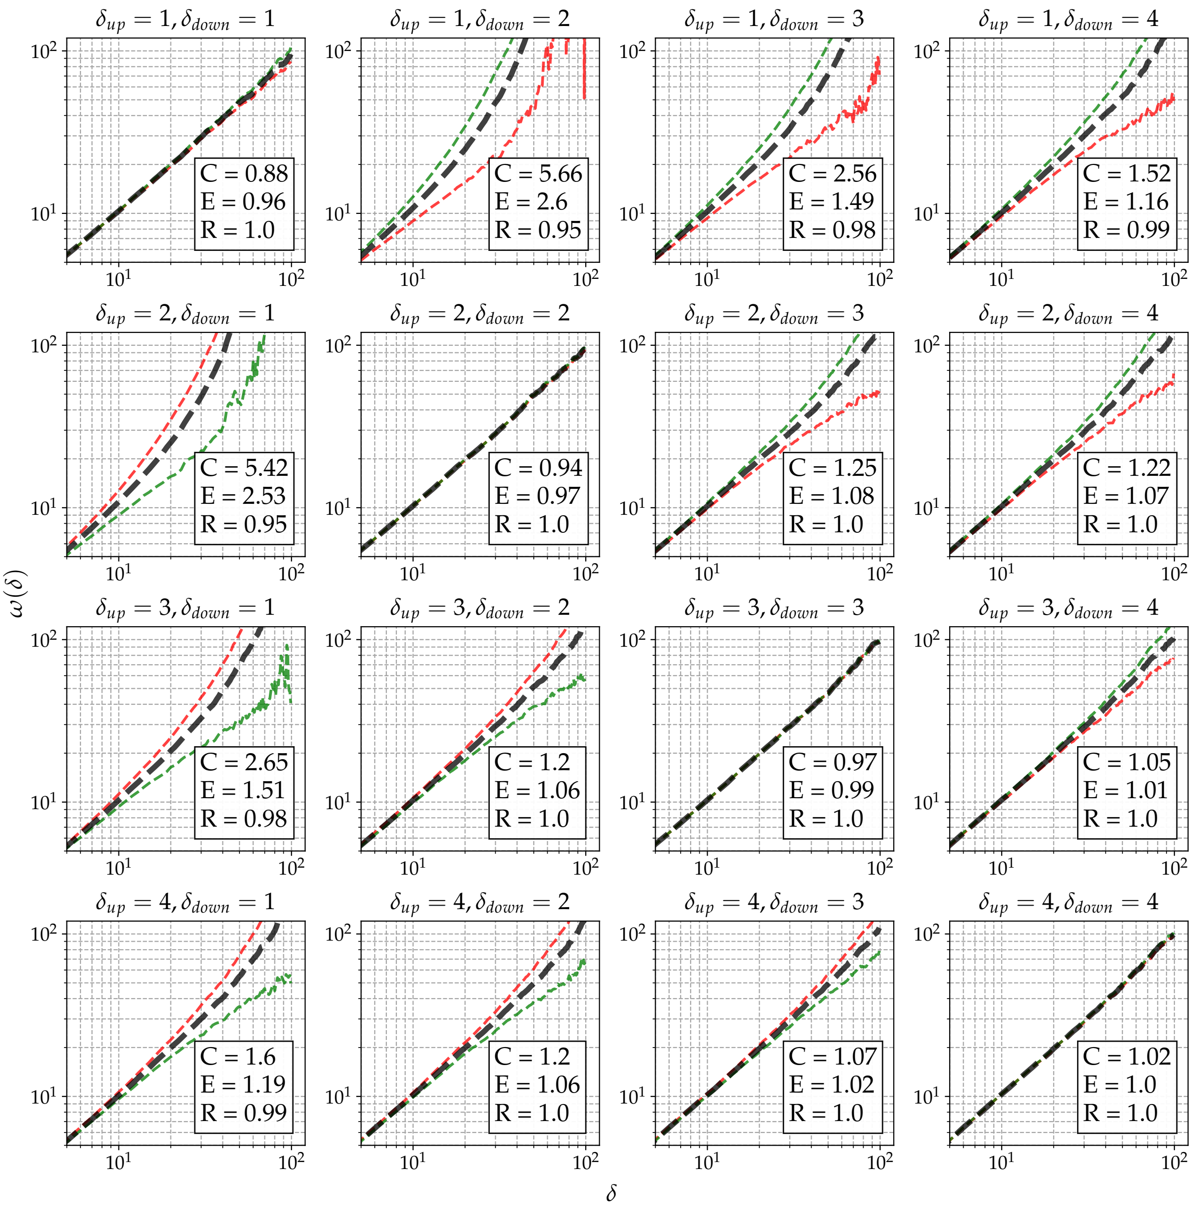

Figure C4. Overshoot scaling laws computed using the time series generated by the entire grid of agents plus one extra trader (see description of figure C3). Green lines correspond to average overshoots computed on the downward trend and red lines to the average overshoots captured on the upward trend. The black line is the average of both lines. Scaling law coefficients C, E and the Pearson product-moment correlation correspond to the average (black) line.

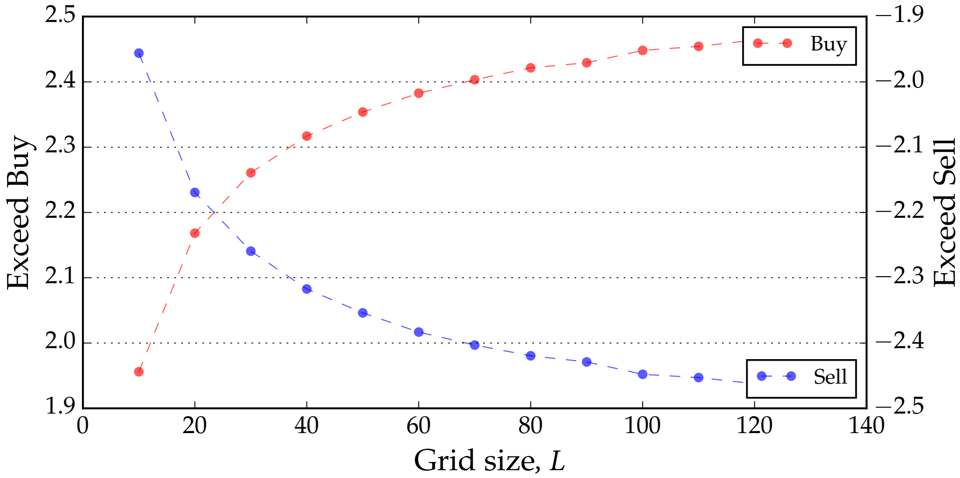

Figure C5. Positive (Exceed Buy) and negative (Exceed Sell) average net volumes as the function of the grid size. Average total net volume is equal to zero.

References

Aloud, M., Fasli, M., Tsang, E., Dupuis, A. and Olsen, R., Modeling the High-Frequency FX Market: An Agent-Based Approach. *Computational Intelligence*, 2017, **33**, 771–825.

Andriani, P. and McKelvey, B., Beyond Gaussian averages: redirecting international business and management research toward extreme events and power laws. *Journal of International Business Studies*, 2007, **38**, 1212–1230.

Antypas, A., Koundouri, P. and Kourogenis, N., Aggregational Gaussianity and barely infinite variance in financial returns. *Journal of Empirical Finance*, 2013, **20**, 102–108.

Arneodo, A., Bouchaud, J.P., Cont, R., Muzy, J.F., Potters, M. and Sornette, D., Comment on"Turbulent cascades in foreign exchange markets". *arXiv preprint cond-mat/9607120*, 1996.

Basu, S., Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis. *The journal of Finance*, 1977, **32**, 663–682.

Bauwens, L., Omrane, W.B. and Giot, P., News announcements, market activity and volatility in the euro/dollar foreign exchange market. *Journal of International Money and Finance*, 2005, **24**, 1108–1125.

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

Blattberg, R.C. and Gonedes, N.J., A comparison of the stable and student distributions as statistical models for stock prices. *The journal of business*, 1974, **47**, 244–280.

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

Bouchaud, J.P., Power laws in economics and finance: some ideas from physics. , 2001.

Bouchaud, J.P., Price impact. *Encyclopedia of quantitative finance*, 2010.

Brown, S.J. and Warner, J.B., Using daily stock returns: The case of event studies. *Journal of financial economics*, 1985, **14**, 3–31.

Chakraborti, A., Toke, I.M., Patriarca, M. and Abergel, F., Econophysics: Empirical facts and agent-based models. *arXiv preprint arXiv:0909.1974*, 2009.

Cheung, Y.W. and Chinn, M.D., Currency traders and exchange rate dynamics: a survey of the US market. *Journal of international money and finance*, 2001, **20**, 439–471.

Chordia, T., Roll, R. and Subrahmanyam, A., Market liquidity and trading activity. *The journal of finance*, 2001, **56**, 501–530.

Christie, A.A., The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of financial Economics*, 1982, **10**, 407–432.

Cohen, L., Diether, K.B. and Malloy, C.J., Supply and demand shifts in the shorting market. *The Journal of Finance*, 2007, **62**, 2061–2096.

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

Cont, R., Potters, M. and Bouchaud, J.P., Scaling in stock market data: stable laws and beyond. *Scale invariance and beyond*, 1997, **7**, 75–85.

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

Danielsson, J. and de Vries, C.G., Tail index and quantile estimation with very high frequency data. *Journal of empirical Finance*, 1997, **4**, 241–257.

Debreu, G., *Theory of value: An axiomatic analysis of economic equilibrium*, 17, 1987, Yale University Press.

Delage, V., Brandlhuber, C., Tuyls, K. and Weiss, G., Multi-Agent based simulation of FOREX exchange market. *Maastricht University, Department of Knowledge Engineering, the Netherlands Teramark Technologies GmbH, Munich, Germany*, 2010.

Di Matteo, T., Multi-scaling in finance. *Quantitative finance*, 2007, **7**, 21–36.

Ederington, L.H. and Lee, J.H., How markets process information: News releases and volatility. *The Journal of Finance*, 1993, **48**, 1161–1191.

Ehrentreich, N., *Agent-based modeling: The Santa Fe Institute artificial stock market model revisited*, Vol. 602, , 2007, Springer Science & Business Media.

Elder, A., *The new trading for a living: psychology, discipline, trading tools and systems, risk control, trade management*, 2014, John Wiley & Sons.

Fama, E.F., Efficient capital markets: A review of theory and empirical work. *The journal of Finance*, 1970, **25**, 383–417.

Farmer, J.D., Lillo, F. *et al.*, On the origin of power-law tails in price fluctuations. *Quantitative Finance*, 2004, **4**, 7–11.

Frank, M.Z. and Sanati, A., How does the stock market absorb shocks?. *Journal of Financial Economics*, 2018, **129**, 136 – 153.

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E., A theory of power-law distributions in financial market fluctuations. *Nature*, 2003, **423**, 267–270.

Gençay, R., Dacorogna, M., Muller, U.A., Pictet, O. and Olsen, R., *An introduction to high-frequency finance*, 2001, Academic press.

Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P., Dodge, Y. *et al.*, Turbulent cascades in foreign exchange markets. *Nature*, 1996, **381**, 767–770.

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

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

Golub, A., Glattfelder, J. and Olsen, R.B., The Alpha Engine: Designing an Automated Trading Algorithm. , 2017.

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

Hasbrouck, J., Measuring the information content of stock trades. *The Journal of Finance*, 1991, **46**, 179–207.

Hausman, J.A., Lo, A.W. and MacKinlay, A.C., An ordered probit analysis of transaction stock prices. *Journal of financial economics*, 1992, **31**, 319–379.

Huberman, G. and Stanzl, W., Price Manipulation and Quasi-Arbitrage. *Econometrica*, 2004, **72**, 1247–1275.

Joulin, A., Lefevre, A., Grunberg, D. and Bouchaud, J.P., Stock price jumps: news and volume play a minor role. *arXiv preprint arXiv:0803.1769*, 2008.

Kaldor, N., Capital accumulation and economic growth. In *The theory of capital*, pp. 177–222, 1961, Springer.

Kempf, A. and Korn, O., Market depth and order size. *Journal of Financial Markets*, 1999, **2**, 29–48.

Kozhan, R. and Salmon, M., The information content of a limit order book: The case of an FX market. *Journal of Financial Markets*, 2012, **15**, 1–28.

Kullmann, L., Töyli, J., Kertesz, J., Kanto, A. and Kaski, K., Characteristic times in stock market indices. *Physica A: Statistical Mechanics and its Applications*, 1999, **269**, 98–110.

Kyle, A.S., Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1985, pp. 1315–1335.

LeBaron, B., Agent-based computational finance. *Handbook of computational economics*, 2006, **2**, 1187–1233.

Lillo, F., Farmer, J.D. and Mantegna, R.N., Econophysics: Master curve for price-impact function. *Nature*, 2003, **421**, 129.

Mandelbrot, B. and Taylor, H.M., On the distribution of stock price differences. *Operations research*, 1967, **15**, 1057–1062.

Mandelbrot, B.B. and Hudson, R.L., *The (mis) behaviour of markets: a fractal view of risk, ruin and reward*, 2010, Profile books.

Mantegna, R.N. and Stanley, H.E., Scaling behaviour in the dynamics of an economic index. *Nature*, 1995, **376**, 46–49.

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

Müller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M. and Morgenegg, C., Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking & Finance*, 1990, **14**, 1189–1208.

Muth, J.F., Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, 1961, pp. 315–335.

Naciri, N. and Tkiouat, M., Economic Agent Based Models. *International Journal of Applied Engineering Research*, 2016, **11**, 5492–5502.

Pagan, A., The econometrics of financial markets. *Journal of empirical finance*, 1996, **3**, 15–102.

Petrov, V., Golub, A. and Olsen, R., Instantaneous Volatility Estimator Based on Directional-Change Intrinsic Time [online]. , 2018. (accessed Nov. 2018).

Samanidou, E., Zschischang, E., Stauffer, D. and Lux, T., Agent-based models of financial markets. *Reports on Progress in Physics*, 2007, **70**, 409.

Schwert, G.W., Why does stock market volatility change over time?. *The journal of finance*, 1989, **44**, 1115–1153.

Scott, L.O., Option pricing when the variance changes randomly: Theory, estimation, and an application. *Journal of Financial and Quantitative analysis*, 1987, **22**, 419–438.

Walras, L., *Elements of pure economics*, 2013, Routledge.

Zhang, L., Mykland, P.A. and Aït-Sahalia, Y., A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 2005, **100**, 1394–1411.

Zhou, B., High-frequency data and volatility in foreign-exchange rates. *Journal of Business & Economic Statistics*, 1996, **14**, 45–52.