



HAL
open science

An agent-based model of intra-day financial markets dynamics

Jacopo Staccioli, Mauro Napoletano

► **To cite this version:**

Jacopo Staccioli, Mauro Napoletano. An agent-based model of intra-day financial markets dynamics. Journal of Economic Behavior and Organization, 2021, 182, pp.331-348. 10.1016/j.jebo.2020.05.018 . halshs-03046657

HAL Id: halshs-03046657

<https://shs.hal.science/halshs-03046657v1>

Submitted on 13 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Staccioli, Jacopo; Napoletano, Mauro

Working Paper

An agent-based model of intra-day financial markets dynamics

LEM Working Paper Series, No. 2018/12

Provided in Cooperation with:

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

Suggested Citation: Staccioli, Jacopo; Napoletano, Mauro (2018) : An agent-based model of intra-day financial markets dynamics, LEM Working Paper Series, No. 2018/12, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/203065>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

INSTITUTE
OF ECONOMICS



Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

An agent-based model of intra-day financial markets dynamics

Jacopo Staccioli †
Mauro Napoletano ††§

† Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy
‡ Observatoire Français des Conjonctures Economiques (OFCE), Sophia Antipolis, France
§ Université Côte d'Azur, SKEMA, CNRS, GREDEG, France

2018/12

June 2018

ISSN(ONLINE) 2284-0400

An agent-based model of intra-day financial markets dynamics*

Jacopo Staccioli^{†‡} Mauro Napoletano^{§†¶}

1st June 2018

Abstract

We propose a parsimonious agent-based model of a financial market at the intra-day time scale that is able to jointly reproduce many of the empirically validated stylised facts. These include properties related to returns (leptokurtosis, absence of linear autocorrelation, volatility clustering), trading volumes (volume clustering, correlation between volume and volatility), and timing of trades (number of price changes, autocorrelation of durations between subsequent trades, heavy tail in their distribution, order-side clustering). With respect to previous contributions we introduce a strict event scheduling borrowed from the Euronext exchange, and an endogenous rule for traders' participation. We find that the latter proves crucial for matching our target stylised facts.

JEL classification: C63, D84, G12.

Keywords: Intraday financial dynamics, Stylized facts, Agent-based artificial stock markets, Market microstructure, High-Frequency Trading.

*We wish to thank Giovanni Dosi and participants to ISCEF 2018 (Paris, France) and EAEPE 2017 (Budapest, Hungary) for helpful comments and insightful suggestions at various stages of this work. The authors acknowledge support from European Union's Horizon 2020 grant No. 640772 – Project Dolfins. The usual disclaimer applies.

[†]Institute of Economics, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33, 56127 Pisa, Italy.

[‡]To whom correspondence should be addressed: ✉ j.staccioli@santannapisa.it

[§]Observatoire Français des Conjonctures Économiques (OFCE), 60 rue Dostoïevski, BP 085 06902, Sophia Antipolis Cedex, France.

[¶]Université Côte d'Azur, SKEMA, CNRS, GREDEG, France.

1 Introduction

During recent years, the availability of advanced technology has been substantially reducing the latency required to operate on financial markets, fostering market activity at increasingly higher frequencies. As opposed to traditional money managers that generally hold their portfolio positions for a lengthy period, ranging from a few days to even months or more regarding so-called ‘value investors’, high-frequency traders aim at reaping a tiny profit for each of a large multitude of buy and sell operations that they execute *within* each trading day, rarely holding their positions overnight. Short-term trading strategies have proved remarkably profitable even during periods of nearly unprecedented financial turmoil (see e.g. Aldridge, 2013). While the trend of progressively shortening the time needed to collect real-time information and post a new order has been in place for many decades, starting with the introduction of high-speed telegraph service and later boosted by the availability of powerful computer systems, a full and agreed understanding of the functioning, potential benefits, and disadvantages of high-frequency trading has yet to be reached. By the same token, many of the statistical properties (stylised facts) that are empirically recognised to be pervasive of intra-day financial market dynamics are still begging for a sound theoretical framework (see Cont, 2011 for an overview).

We propose a parsimonious agent-based model of a financial market at the intra-day time scale that is able to *jointly* reproduce many of the empirically validated stylised facts. These include properties related to returns (leptokurtosis, absence of linear autocorrelation, volatility clustering), trading volumes (volume clustering, correlation between volume and volatility), and timing of trades (number of price changes, autocorrelation of durations between subsequent trades, heavy tail in the distribution of such durations, order-side clustering).

In the last few decades, the still flourishing literature on agent-based models (see Hommes (2006) for a general review, and in particular LeBaron (2006) for applications to financial markets) has proved invaluable in investigating and replicating the emergence of stylised facts that are hardly reconcilable with the representative agent paradigm.¹ However, the vast majority of the proposed models (some of the milestones include: Arthur et al., 1997; Levy et al., 1994; Lux, 1995, 1998; Lux and Marchesi, 2000) typically focus uniquely on a subset of the whole ensemble of recognised stylised facts, mainly consisting of properties that are time-scale invariant. These generally include properties related to rates of return, such as leptokurtosis, absence of linear autocorrelation, and volatility clustering. Other stylised facts concerning the timing of orders posting and trades execution are often neglected. This is partially due to the acknowledged difficulty of defining a reasonable mapping from the iterations of a computer simulation to proper calendar time (see e.g. Cioffi-Revilla, 2002). A notable exception is Kluger and McBride (2011), who propose a model that replicates the intra-day U-shaped seasonality in market activity, i.e. the tend-

¹Some of these models (often dubbed ‘heterogeneous agent models’) are low-dimensional and mathematically tractable; others, are too complex to be investigated analytically and rely on extensive numerical simulations. Our work is intended to contribute to the latter strand of literature.

ency of exchanged volumes to peak during the early morning just after market opening and late afternoon just short of market closing, leaving a trough around lunch-time.

To the best of our knowledge, no previous study has ever addressed the simultaneous emergence of all the stylised facts that shape financial dynamics at the intra-day level. We therefore attempt at filling this gap by proposing a methodological solution to the time mapping problem and by identifying a set of minimal building blocks whose combination gives a model that is able to reproduce many of the solicited facts.

Our model relies on three main ingredients. The first consists of a behavioural specification on behalf of the traders, which is typical of many previous models in the literature. We assume that traders are of two types: fundamentalists and chartists. Fundamentalists only take into account the fundamental value of the security (which we shall assume constant and common knowledge), buying the asset if undervalued and selling it if overvalued. Chartists conversely only take into account the recent history of price changes, extrapolating the underlying trend if they are followers or counteracting it if they are contrarians. This specification is justified by empirical surveys of financial practitioners' behaviour (see e.g. Frankel and Froot, 1990). The second ingredient, which to our knowledge has never appeared in any previous contribution, amounts to a well-defined scheduling of events. We borrow the exact time structure of a trading day on a real financial market, namely the EURONEXT, and we design our simulations according to the sequence and durations of the different phases therein (see Euronext, 2017). The latter consists, in chronological order, of a morning order accumulation phase, an opening batch auction, a lengthy phase of real-time order matching according to a continuous double auction, a pre-closing order accumulation phase, and a closing batch auction. We believe that imposing a strict and realistic schedule on the unfolding of events enables to devise a sound and plausible correspondence between simulation ticks (which we shall identify with seconds) and calendar time. Microstructure details about the central order book are also compliant with EURONEXT specifications. The last ingredient of the model is an endogenous mechanism for traders participation. We assume that traders (of either type) are more willing to engage in trading, the higher the price change (of either sign) realised in the immediate past. The idea is that large realised (absolute) returns signal the possibility of reaping further profit in the future. Note that in the following we shall not impose any short-sale restriction. A similar scheme is devised in Ghoulmie et al. (2005) and Jacob Leal et al. (2016). In spite of being extremely simplistic, we find that this activation mechanism proves crucial for matching our target stylised facts, specifically those related to the timing of trades execution. An additional conceivable ingredient, commonly adopted in financial models akin to ours, is a switching scheme between the fundamental and chartist strategies. In many contributions this is known to foster volatility clustering (see e.g. Kirman and Teyssière, 2002; Lux and Marchesi, 2000, and for a discussion Cont, 2007). We find however that this component is irrelevant in our setting, in which volatility clustering arises purely from the interaction of heterogeneous traders and is especially influenced by the trend-following momentum on behalf of chartists.

The next Section provides an overview of the various stylised facts that characterise high frequency financial dynamics. Section 3 describes in detail the various assumptions of our

model. [Section 4](#) reports the results of numerical simulations under different scenarios. Finally, [Section 5](#) concludes.

2 Stylised facts

Here we describe the various statistical properties that characterise the intraday dynamics of many financial markets and that our model aims at reproducing. Some of them are recognised to apply across different time scales while others require a proper intraday setting to be analysed (we refer to Cont, [2001](#), [2011](#) for a more detailed account). While the former are mainly related to the statistical properties of returns and have been already studied and replicated in a number of agent-based models lacking of a rigorous definition of calendar time, the latter require a more explicit architecture in terms of microstructure and constitute the main contribution of our work. We list and discuss them in order.

SF1 – Leptokurtic returns The unconditional distribution of returns is characterised by a heavier tail with respect to the Gaussian distribution (Fama, [1965](#); Kon, [1984](#)). The magnitude of excess kurtosis is typically inversely related to the time scale of analysis. This finding stands at sharp odds with the normality assumption adopted in a number of models, most notably the Black-Scholes formula.

SF2 – Absence of autocorrelation of (raw) returns The time series of (raw) rates of return exhibits a statistically significant serial correlation for a very short amount of time, quickly decaying to zero afterwards. Intuitively, should there be more predictable autocorrelation structure, this information could be used to perform ‘*statistical arbitrage*’ with positive profits (Mandelbrot, [1971](#)).

SF3 – Volatility clustering While the linear autocorrelation of returns displays very little structure, the autocorrelation of non-linear functions such as the absolute value or the squared value of returns is usually positive and tends to decay at a much slower pace. Therefore, while the signs of future returns are not readily predictable, their magnitudes are, and tend to cluster in time, giving rise to prolonged periods of low volatility followed by periods of high volatility (Andersen and Bollerslev, [1997](#); Mandelbrot, [1963](#)). This clearly suggests that the series of returns is not independent.

SF4 – Leverage effect The leverage effect or asymmetric volatility (Black, [1976](#)) captures the asymmetric tendency of volatility to be higher during price drops rather than during price surges. This translates into the negative correlation between price volatility – e.g. absolute returns – and the (raw) returns of the asset (Aït-Sahalia et al., [2013](#); Bollerslev et al., [2006](#); Bouchaud et al., [2001](#)).

SF5 – Number of price changes per day Under a cross section perspective, the number of effective trades per day is clearly related to the degree of liquidity of the market and

is typically linked to the capitalisation of the underlying security. Over time, moreover, there is a tendency of reduction in the time needed to execute a market order, fostering the submission of an increasingly larger number of orders, eventually leading to an increasing frequency of actual trades. Nowadays, for blue-chips in highly liquid markets and in the absence of 'disruptive' fundamental news, this number is often around 10,000, with a substantial degree of variance (Bonanno et al., 2000; McNish and Wood, 1991).

SF6 – Autocorrelation of durations between subsequent trades Within continuous double auctions, the actual timing of transactions is endogenous since a freshly submitted order might not find a compatible crossing order already stored in the book. Therefore, the time intervals between subsequent transactions is both random and tightly linked to the previous history of orders posting. Empirically, these durations display positive autocorrelation – translating in clustered periods of frequent transactions followed by periods of sporadic transactions.

SF7 – Fat-tailed distribution of durations between subsequent trades The distribution of the durations defined in SF6 reveals a heavier tail with respect to an exponential distribution, that would be instead expected if traders submitted their orders in a non-correlated timely fashion (Raberto et al., 2002).

SF8 – Order-flow clustering The arrival of orders over time to the central order book is clustered with respect to the side of intended transaction: buy orders are more likely to follow previous buy orders, while sell orders are more likely to follow sell orders (Biais et al., 1995).

SF9 – Autocorrelation of volumes The quantities exchanged during successive trades display significant positive serial correlation (Campbell et al., 1993; Engle, 2000; Gallant et al., 1992). This is true across different time aggregation units and both for indices and individual stocks.

SF10 – Correlation between volumes and volatility Price variability and trading volumes display positive correlation (Foster, 1995; Tauchen and Pitts, 1983). The underlying idea is that the flow of information acts as a common determinant of both changes in prices and traded quantities.

SF11 – U-shaped activity Market activity throughout the day displays a strong seasonality, with peaks of exchanged quantities in the early morning after market opening and in the late afternoon in the vicinity of market closing, and a relative more tranquil period in the hours around lunch-time (Jain and Joh, 1988; Lockwood and Linn, 1990).

In what follows, we aim at developing a simple and parsimonious model which is nonetheless capable of *jointly* reproducing all the aforementioned stylised facts, with the ex-

ception of the intraday volume seasonality² (SF11), which is unobtainable by construction in our setting, as will be clear later, and the leverage effect (SF4), for which we believe a more complex behavioural specification is needed.

3 The model

Consider an order-driven financial market in which a single long-lived stock is traded by a population of heterogeneous agents. In line with the empirical literature on practitioners' behaviour in financial markets pioneered, among others, by Frankel and Froot (1990), Allen and Taylor (1990), Taylor and Allen (1992), and more recently by Menkhoff (2010), we consider two trading strategies: fundamentalist and chartist. A fundamentalist trader believes that the price of a security will quickly revert to its fundamental value; a chartist (or technical) trader, instead, believes that the future price of a security can be predicted using the trend of past realised market outcomes. Since we are interested in modelling short-term dynamics, we assume that the security pays no dividend and there is no 'fundamental' news circulating during this time span. In this sense, besides an additive i.i.d. noise component incorporated in both strategies whose only purpose is to ensure that trading doesn't jam, the dynamics of prices and returns is endogenously determined by the interaction of the two strategies, and realised volatility is actually excess volatility.

3.1 Timing and market setting

Since we are interested in describing the dynamics of a generic stock at a well-defined time scale – the intra-day level – we need to devise a mechanism that maps the iterations of our agent-based model to proper calendar time. This is a notoriously daunting and controversial task within the agent-based literature (see e.g. Cioffi-Revilla, 2002). To address this issue, we impose a strict global schedule to the sequence of events. In particular, we design our simulations to closely replicate the timing structure of an existing stock market, namely the EURONEXT. A typical trading day on the EURONEXT exchange unfolds as follows:

at 7:15am the trading day starts with the pre-opening phase in which orders accumulate on the central order book without any transactions taking place;

at 9:00am a (batch) opening auction takes place, matching the orders submitted during the pre-opening phase and determining the opening price;

from 9:00am to 5:30pm the market operates according to a continuous double auction, and the introduction of a new order immediately generates one or more transactions if there are matching orders on the opposite side of the book. This phase is dubbed the 'main trading session';

²Kluger and McBride (2011) provide an agent-based model that reproduces the U-shaped nature of intraday volumes, although they don't discuss the whole ensemble of the stylised facts listed above.

at 5:30pm pre-closing phase starts, in which matching of orders is discontinued and, as in the pre-opening phase, orders accumulate with no transaction taking place;

at 5:35pm the closing auction takes place, matching the orders submitted during the pre-closing phase and determining the closing price of the day.

from 5:35pm to 5:40pm orders can be entered for execution at the closing price only. This phase is dubbed ‘trading at last’.

With the exception of the trading at last phase³, we model our trading day according to the schedule above, and we identify a single iteration of the model with a calendar second. Hence, the pre-market phase corresponds to 6,300 time steps (1 hour and 45 minutes), the main trading session to 30,600 time steps (8½ hours), and the pre-closing phase to 300 time steps (5 minutes). A whole trading day consists of 37,200 simulation steps of our model.

At every time step some of the traders are activated, they proceed to form their expectations about future performance of the security, and submit limit orders accordingly. When an order is submitted, it is either stored on the central order book or matched (if possible), depending on the current phase of the trading day. If matched, the order gives rise to one or more trades, the relevant quantities are exchanged, and a new price is disseminated. The central order book follows the usual price-time priority rule.

We devise two alternative mechanisms for traders activation, one exogenous and one endogenous. In the first, a single trader is activated at each time step, randomly selected among the population. This activation scheme is similar to the one employed by Chiarella and Iori (2002). In the second, we follow Jacob Leal et al. (2016) and we assume that traders’ activation is endogenous in the following sense: at every time step all traders decide whether they are willing to submit an order by comparing the last recorded price change (in absolute value) to a trader-specific and time-varying threshold, drawn from a common distribution with positive support. In particular, trader i is active at time t if $|r_\tau| > \delta_{i,t} \sim |\mathcal{N}(0, \sigma_\delta^2)|$, where $\tau < t$ denotes the last time in which a trade occurred. If multiple agents are active at time t , they engage in trading in randomised order. If no trader is endogenously activated at time t , then with a certain probability $\phi > 0$ the mechanism falls back to the baseline activation scheme, and a randomly selected trader is asked to submit an order. While the first mechanism is useful as a baseline scenario to describe and test the functioning of the model, we discover that the second mechanism is better suited for replicating our target stylised facts. In particular, this scenario admits both crowded and uneventful periods in which either many or no orders are submitted, and contribute to clustering of volumes and of the order-flow. Endogenous activation schemes are known to be crucial for matching realistic financial dynamics. Pellizzari and Westerhoff (2009) devise a similar rule, based on past profits.

³We don’t model this phase since by construction has no influence on the price of the security, and is therefore deemed irrelevant with respect to our objective stylised facts.

3.2 Traders' behaviour

Traders form expectations about the future (log) return over a certain time horizon h as follows:

$$\hat{r}_{i,t+h}^F = w_i^F \cdot \log\left(\frac{p^F}{p_t}\right) + \varepsilon_t \quad (3.1)$$

$$\hat{r}_{i,t+h}^C = w_i^C \cdot \log\left(\frac{p_t}{p_{t-h}}\right) + \varepsilon_t \quad (3.2)$$

The superscript F (respectively, C) identifies the fundamentalist (respectively, chartist) strategy, $w_i^F \sim |\mathcal{N}(0, \sigma_F^2)|$ and $w_i^C \sim \mathcal{N}(\mu_C, \sigma_C^2)$ are trader-specific coefficients that describe how 'aggressively' the underlying strategy is implemented, $p^F > 0$ denotes the fundamental price of the security, $h \in \mathbb{N}_+$ measures the horizon the trader operates within, and $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is a common i.i.d. noise component. The fundamental price p^F in eq. (3.1) is common knowledge of all traders and weight w_i^F quantifies how quickly the price of the stock is expected to revert to its fundamental value. From eq. (3.2), it is clear that all chartists use only the last realised return over the time-span h to form their expectation. Weight w_i^C measures the extent to which traders believe the future return over period h will match its past figure. This assumption – that helps in containing the dimensionality of the model⁴ – stands at variance with previous literature which usually assume a weighted moving average (typically exponentially or linearly) over *multiple* past returns. However, given our intra-day setting, we believe that the short memory of chartists mimics more closely the fast response of high-frequency traders to suddenly realised signals. Notice also that we admit an imbalance between trend followers and contrarians, depending on the value of the mean μ_C of the distribution of chartists' weights w_i^C .

Once a trader has formed her expectation about the future return, she submits a limit order to the central order book. A limit order is a triple {price, quantity, validity} such that: price equals the expected prevailing price at the end of period $t + h$, rounded to the nearest tick; quantity is always fixed to one unit, carrying a positive (respectively, negative) sign if the order is to be stored on the buy (respectively, sell) side of the central order book, depending on whether the trader expects the price to increase or decrease; validity, namely the time after which the order expires and is automatically deleted from the central order book, is set to equal the horizon of the expectation. We assume that all traders have unlimited access to external credit at a zero interest rate, so that they can either short-sell or leverage-buy the stock without bound. In other terms, traders don't face a budget constraint; nevertheless, they are prevented from borrowing an infinite amount of money by the unitary quantity rule. To sum up, a limit order $\ell_{i,t}$ submitted by trader i (either fundamentalist or chartist) at time t takes the form:

$$\ell_{i,t} = \{\text{round}(p_t \cdot \exp(\hat{r}_{i,t+h}), \text{tick}), \text{sign}(\hat{r}_{i,t+h}), t + h\} \quad (3.3)$$

⁴The gains in terms of parsimony are due to the fact that we don't need to quantify the memory of the traders (or even worse, a distribution thereof), and a rate of decay of the importance of remote past history.

where $\text{round}(\cdot)$ denotes the rounding function, $tick$ is the minimum price increment/decrement (a parameter of the market), and $\text{sign}(\cdot)$ is the sign function.

While we don't include the early cancellation of an order within the strategy set of a trader, we devise the following automatic cancellation rule: when a trader submits a new order, all other orders already submitted by the same trader and stored on the book that are inconsistent with the new expectation are automatically cancelled. These include all orders stored on the opposite side of the book and those orders whose underlying price is deemed unfavourable given current expectations. For example, when a buy (respectively, sell) order is issued at price \tilde{p} , all sell (respectively, buy) orders, and all buy (respectively, sell) orders whose price is greater (respectively, less) than \tilde{p} , are automatically cancelled. The first condition ensures that a trader never trades with herself, i.e. it rules out the possibility that two orders submitted by the same trader are matched together. The second condition ensures that in case a trader is currently willing to buy (sell) the security at a certain price, she is no longer willing to buy (sell) at a higher (lower) price, as per orders submitted under possibly different beliefs.

It is important to note that no reference whatsoever to any specific time of the day appears in either eq. (3.1) or eq. (3.2). In other words, none of the traders knows 'what time it is' when asked to submit an order, and behaves identically throughout every phase and instant of the trading day. This implies that, by construction, our model is unlikely to validate SF11, and that any spike in market activity observed in our series has the same probability of occurring during morning, lunch, or afternoon time.

4 Numerical simulations

In spite of the very simple behavioural rules that we assume, the complexity associated with the endogenous nature of a limit order book dynamics prevents us from studying the system analytically and come up with a closed form solution. We thus follow the standard practice, typical of agent-based models, of numerically simulating the system and then performing the relevant statistical tests on the generated series.

We start by fixing a few parameters and design principles that are kept stable across our simulations. The market is populated by $N = 1,000$ traders; the fundamental price of the stock is constant and equals $p^F = 100$, while the $tick$ value – i.e. the smallest possible increment or decrement of the price – equals 0.001. At the beginning of the simulation the price is set to equal its fundamental value, i.e. $p_0 = p^F$, and all chartists are provided with a history of past prices between $t = -h$ and $t = 0$ that evolves (backwards) as a pure random walk whose increments are given by the same noise component ε_t present in eqs. (3.1) and (3.2). Finally, we fix the horizon of traders' expectations h to 1,000 seconds (simulation time steps); incidentally, this value equals the expected duration between two consecutive activations of a same trader within the exogenous activation scheme, given the number of traders N .

At the finest level of granularity, our simulations yield time series of the relevant quantities that are *irregular* since, by construction, trade emerges endogenously when at least

Parameter	Value
ε_t	$\mathcal{N}(0, 5e-5)$
w_i^F	0
w_i^C	0
δ_t	$+\infty$
ϕ	1

Table 1: Parameters value and initial conditions for the NT scenario.

two crossing orders are stored on the central order book. In order to perform the statistical analysis needed to validate our model against the stylised facts listed in [Section 2](#), we rescale the relevant time series by pooling the stream of trade messages into regular time windows of one calendar minute each. The minute-by-minute price (respectively, volume) series corresponds to the average (respectively, sum) of the underlying trading prices (volume) during that minute. Following [Section 3.1](#), the main trading session consists of 510 minutes.

We proceed by simulating the model numerically⁵ under three different scenarios, following a bottom-up approach. In [Section 4.1](#) we only include purely noise traders; in [Section 4.2](#) we investigate the interaction between fundamentalists and chartists under the baseline exogenous activation; in [Section 4.3](#) we discuss the effect of introducing our endogenous activation scheme. Finally, in [Section 4.4](#) we perform some complementary sensitivity analysis. The results that we show correspond to averages across 100 Monte-Carlo simulations of a fully fledged trading day (see [Section 3.1](#)). All confidence intervals are set at the 95% level.

4.1 Noise traders only

This scenario, which we dub NT, is useful to properly disentangle the dynamic properties implied by market microstructure details from those implied by our assumptions about specific trader behaviour. Noise traders do not condition their investment on any market-related variable; rather, they “*trade on noise as if it were information*” (Black, 1986). Given our formulation, we set all the w_i ’s in eqs. (3.1) and (3.2) to zero, such that the expected return for each trader will only depend on the i.i.d. noise component ε_t .

Table 1 summarises the specific parametrisation. By setting δ_t to infinity we rule out endogenous activation, and by setting $\phi = 1$ we ensure that exactly one trader is activated at every time step t .

⁵The model is coded in C++11 and largely exploits the object-oriented programming paradigm, defining *classes* for traders, for the central order book, and for the order data structure. The code supports the execution of fully parallel Monte Carlo simulations, using the OpenMP framework. Random number generation relies on the 32-bit Mersenne Twister, as implemented in the C++ Standard Library (`std::mt19937`). Parameters and initialisation for all the Monte Carlo simulations are passed through a single json file during run-time, so that the code needs not be (re)compiled every time a new scenario is simulated. The file is parsed using the `jsoncpp` library. Each Monte Carlo simulation returns a SQLite database file containing the associated initialisation and a stream of messages from the central order book, each corresponding to a successful transaction (each message reports the current POSIX timestamp, bid, ask, transaction price, quantity, and depth of the book for both sides). The output databases are then imported and analysed using R.




Figure 1: Main stylised facts under the noise traders scenario. (a): evolution of price of a typical trading day (the dashed line denotes the fundamental price); (b): returns of a typical trading day (same as for plot (a)); (c): autocorrelation function of (raw) returns; (d): autocorrelation function of absolute returns; (e): autocorrelation function of inter-trade times; (f): quantiles of inter-trade times vs. theoretical quantiles of an exponential distribution; (g): autocorrelation function of volumes; (h): autocorrelation function of central order book imbalance; (i): scatter plot of volume and volatility with linear regression line; (j): boxplot of correlation between returns and leading squared returns (solid line indicates zero).

Fig. 1 pictures the relevant plots under this scenario. Panel (a) shows the evolution of the minute-by-minute market price for a typical trading day while panel (b) reports its increments. The average number of trades per day under this scenario is 14865. If microstructure effects were completely irrelevant, given our limit price function (3.3), the time series of realised returns would share the same statistical properties of the i.i.d. series of expected returns. We find that the Ljung-Box statistic strongly rejects (p -value < 0.001) the null hypothesis of independence. This is also visible in panel (c) which shows the autocorrelation function of returns; evidently, positive autocorrelation for the first lag is substantial, and for the second lag is very close to the confidence threshold. The Augmented Dickey-Fuller (ADF) test doesn't reject (p -value < 0.001) the presence of a unit root within the price series. Prices are therefore well approximated by a random walk, although its increments are not independent. Moreover, the kurtosis of the sample distribution of returns, $\kappa \approx 3.42$, is only negligibly higher than that of expected returns, that by construction equals 3. We conclude that our microstructure setup does force a time dependence character into the resulting series, although it lasts for just a couple of minutes. Panel (d) pictures the autocorrelation function of the absolute value of returns. Its rate of decay is very high and only the first lag is significant; we conclude that volatility clustering is not present. Panels (e) and (f) relate to the properties of time durations between subsequent trades. Panel (e) shows the autocorrelation function of such durations, while panel (f) pictures a quantile-quantile plot of their distribution, compared to a fitted exponential distribution. The autocorrelation function is negative for the first few lags, and the distribution has a tail that is thinner than that of an exponential distribution. The following two panels suggest that there is no correlation structure in either the exchanged volumes of the asset (panel (g)), nor in the clustering of buy and sell orders stored in the book (panel (h)). Panel (i) reports a negative relationship between exchanged volumes and volatility, instead of the predicted positive correlation. Finally, panel (j) suggests that any leverage effect is absent in our series.

The first column on the right hand side of Table 4 compares these results with our objective stylised facts. It is evident that we need to impose some more structure on the behaviour of the traders in order to obtain a more realistic dynamics.

4.2 Fundamentalists and chartists

In this scenario (FC), we move a step forward by including our fundamentalist and chartist specifications, following eqs. (3.1) and (3.2). On the one hand, fundamentalist traders anchor the price dynamics to a neighbourhood of the fundamental price p^F ; on the other hand, chartists tend to exacerbate or counteract the prevailing trend, depending on their being trend followers or contrarians. The overall stability of the system depends on the balance between fundamentalists and chartists and, within the latter group, between followers and contrarians. Intuitively, the stronger the magnitude of trend following behaviour, the wider the divergence of price from p^F , either upwards or downwards. The parametrisation we propose, reported in Table 2, yields a mean-stationary price series (ADF test p -value < 0.001), and therefore a stable dynamics. Note that we set the value of $\mu_C > 0$.

Parameter	Value
ε_t	$\mathcal{N}(0, 5e-5)$
w_i^F	$ \mathcal{N}(0, 0.001) $
w_i^C	$\mathcal{N}(0.01, 0.1)$
δ_t	$+\infty$
ϕ	1

Table 2: Parameters value and initial conditions for the FC scenario.

Reasonably, the overall sentiment among the crowd of chartists predict a self-reinforcing dynamics, rather than self-opposing. Moreover, as will become clear later in [Section 4.4](#), we find that this assumption fosters the fat-tailedness character of inter-trades durations, and thus helps in replicating our objective stylised facts.

Fig. 2 pictures the relevant plots under this scenario. As expected, the evolution of the price series is more ‘centered’ around the fundamental value p^F with respect to the NT scenario thanks to the fundamentalists’ anchoring behaviour (panel (a)). Moreover, the presence of chartists introduces a persistence character in the dynamics of returns: the presence of both trend followers and contrarians is crucial because their effect on the autocorrelation function of returns cumulates in absolute value, but cancels out when the sign is taken into account. This is clearly visible in panels (c) and (d).⁶ Intuitively, while we allow a slight imbalance between followers and contrarians, a larger imbalance would have the effect of adding memory to the autocorrelation function of (raw) returns, which is contradicted by empirical evidence. The average number of trades, 14798, is in line with the previous scenario. The kurtosis of minute returns has instead increased to 21.3. As predicted in SF1, the kurtosis decreases with the time window and reverts back to 3, i.e. to statistical *normality*, for 15-minute returns. Panels (e), (f), (g), (h), (i), and (j) are qualitatively similar to the NT case. The timing structure of orders submission and matching is not substantially influenced by the presence of the new behavioural specification; exchanged volumes display no persistence character either.

The second column on the right hand side of Table 4 suggests that there is an improvement with respect to the noise traders scenario: volatility clustering and leptokurtosis of price returns are correctly matched after introducing our fundamentalist and chartist strategies.

4.3 Endogenous activation

In this scenario, which we call EA, we add the last ingredient to the mix of fundamentalists and chartists, namely the endogenous activation scheme outlined in [Section 3.1](#). The ultimate goal is to retain the good properties encountered in the FC scenario, namely those pertaining to the distribution and dependence of returns, and improve those related to the timing and volume of trade.

⁶In a separate experiment (not shown) we set $w_i^C \sim \pm |\mathcal{N}(\mu_C, \sigma_C^2)|$, i.e. we include either trend-followers or contrarians but not both. In this case we find that the autocorrelation function of returns (panel (c)) and of absolute returns (panel (d)) look very similar and thus fail to validate our target stylised facts SF2 and SF3.

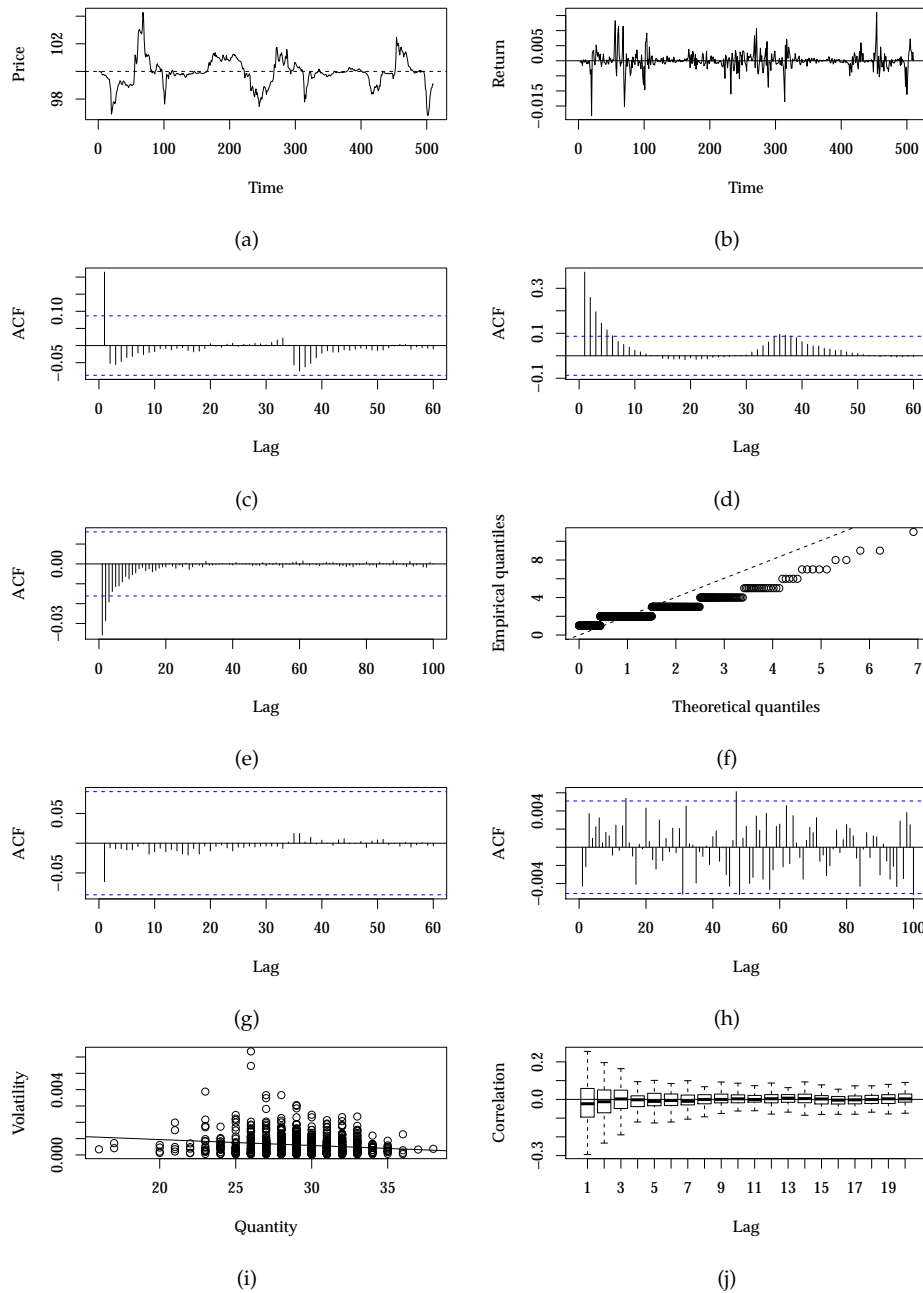


Figure 2: Main stylised facts under the fundamentalists vs. chartists scenario. (a): evolution of price of a typical trading day (the dashed line denotes the fundamental price); (b): returns of a typical trading day (same as for plot (a)); (c): autocorrelation function of (raw) returns; (d): autocorrelation function of absolute returns; (e): autocorrelation function of inter-trade times; (f): quantiles of inter-trade times vs. theoretical quantiles of an exponential distribution; (g): autocorrelation function of volumes; (h): autocorrelation function of central order book imbalance; (i): scatter plot of volume and volatility with linear regression line; (j): boxplot of correlation between returns and leading squared returns (solid line indicates zero).

Parameter	Value
ε_t	$\mathcal{N}(0, 5e-5)$
w_i^F	$ \mathcal{N}(0, 0.001) $
w_i^C	$\mathcal{N}(0.01, 0.1)$
δ_t	$ \mathcal{N}(0, 0.3) $
ϕ	$1/3$

Table 3: Parameters value and initial conditions for the EA scenario.

The endogenous activation scheme aims at capturing the ever more common high-frequency nature of financial markets. A crowd of traders, many of which are algorithmic machines, typically respond very quickly to a newly posted signal and engage in trading with and against each other for a while until they coordinate on a new price⁷. This creates a signature in the series of trading times, in which periods of no trade, sometimes lasting for several tens of seconds, are followed by the submission of a multitude of orders to the central order book within a matter of seconds, giving rise to a substantial number of trades.

Table 3 summarises the specific parametrisation that we employ. The value of $\sigma_\delta^2 = 0.3$ is such that, on average, exactly one trader is endogenously activated at time t in response to a realised absolute return $|r_{t-1}| \approx 0.000375$, whereas the average absolute return in the FC scenario is approximately 0.0003. This means that most of the time traders are not endogenously activated, and the fallback exogenous activation scheme takes over, with probability $\phi = 1/3$. However, due to the leptokurtic nature of returns (SF1), there exist periods in which a much larger-than-average price change takes place, and a multitude of traders are willing to submit new orders at the same time. Moreover, the price change generated by such turbulent event is likely to be itself larger than the δ_t threshold for a number of traders, possibly triggering a new wave of crowded endogenous activation in the next period, ultimately lengthening the duration of the adjustment.

Fig. 3 pictures the relevant plots under this scenario. Panels (a) to (d) are qualitatively similar to those of scenario FC, suggesting that the good properties of returns generated in the latter setting have not been compromised by the new activation assumption. Leptokurtosis has increased to a minute-by-minute figure in excess of 200, decreasing to around 22 for 15-minute returns, and 8.5 for 30-minute returns. The average number of trades, 11739, has decreased as a result of the new activation scheme, but is still a perfectly acceptable level for liquid traded securities (Cont, 2011). The main benefits of the endogenous activation scheme are noticeable in the subsequent panels of Fig. 3. For the first time, panel (e) shows a strong and very slowly decaying autocorrelation in inter-trade durations (SF6), and the quantile-quantile plot in panel (f) suggests that the tail of their distribution is fatter than exponential (SF7). Moreover, both volumes (panel (g)) and order-flow (panel (h)) are

⁷In principle, such a signal can arise either from *within* the order book, e.g. as a disruptive newly submitted order, or from *outside*, in which case it is related to fundamental news about the asset. Empirically, it has been shown that only a fraction of realised volatility is attributable to freshly available news about dividends, prospective earnings, or other crucial balance sheet and macroeconomic variables (see e.g. Cutler et al., 1989; Shiller, 1981). In our model no news is ever released and all traders agree on a constant fundamental value; thus, all the signals come from within the order book, and are the result of sheer trading activity by the traders. The totality of the generated volatility is *excess volatility*.

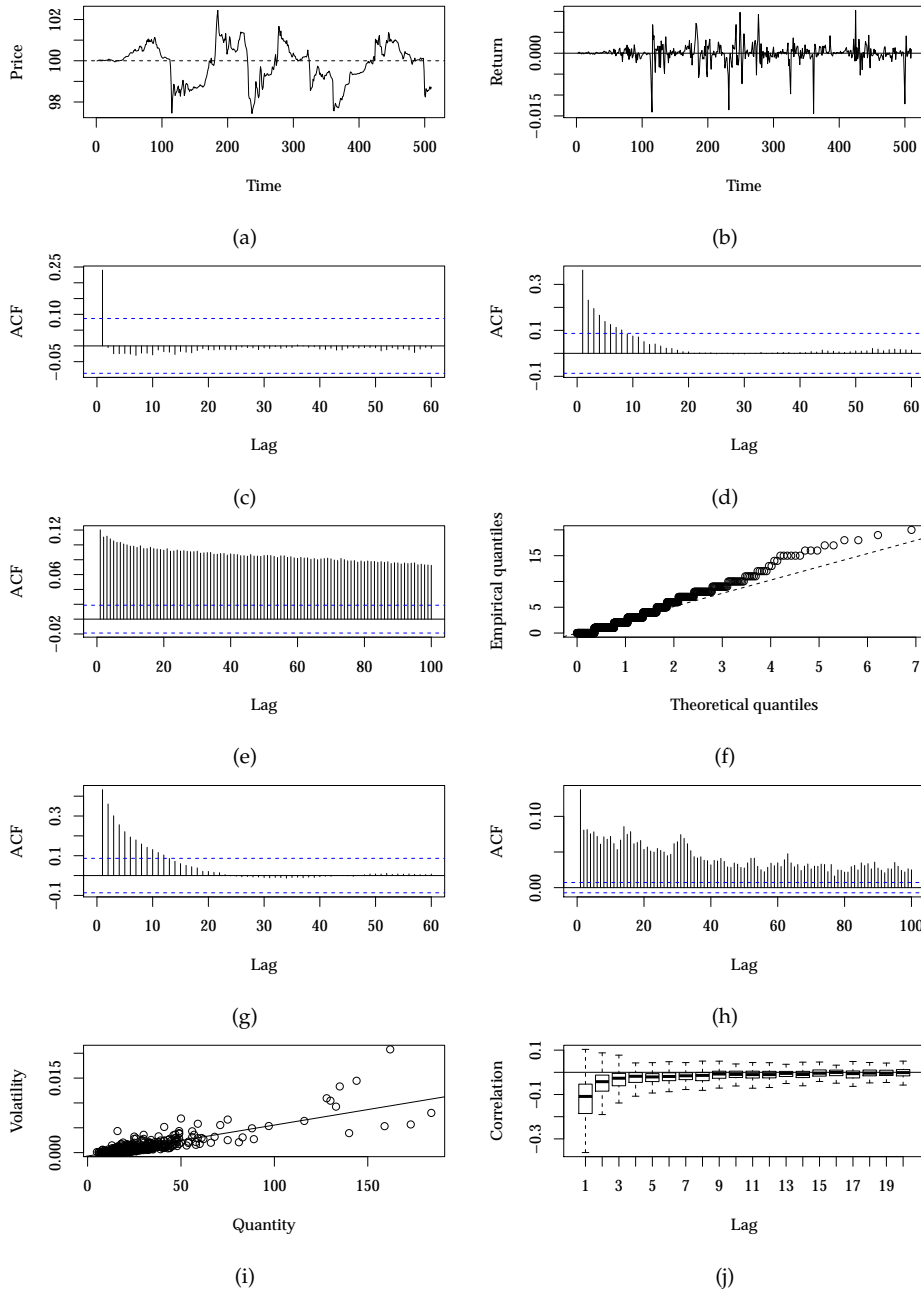


Figure 3: Main stylised facts under the fundamentalists vs. chartists scenario with endogenous activation. (a): evolution of price of a typical trading day (the dashed line denotes the fundamental price); (b): returns of a typical trading day (same as for plot (a)); (c): autocorrelation function of (raw) returns; (d): autocorrelation function of absolute returns; (e): autocorrelation function of inter-trade times; (f): quantiles of inter-trade times vs. theoretical quantiles of an exponential distribution; (g): autocorrelation function of volumes; (h): autocorrelation function of central order book imbalance; (i): scatter plot of volume and volatility with linear regression line; (j): boxplot of correlation between returns and leading squared returns (solid line indicates zero).

stylised fact		scenario		
		NT	FC	EA
SF1	leptokurtic returns	X	✓	✓
SF2	no linear autocorr.	✓	✓	✓
SF3	volatility clustering	X	✓	✓
SF4	leverage effect	X	X	X
SF5	# price changes	✓	✓	✓
SF6	autocorr. durations	X	X	✓
SF7	fat-tailed durations	X	X	✓
SF8	order-flow clustering	X	X	✓
SF9	autocorr. volumes	X	X	✓
SF10	volume/volatility corr.	X	X	✓
SF11	U-shaped activity	X	X	X

Table 4: Replication of the stylised facts within all the simulated scenarios.

clustered (matching respectively SF9 and SF8). Panel (i) shows a positive and significant relationship between volumes and volatility (p -value < 0.001) (as per SF10). An analogously significant relationship holds also for pooled series at 15-minute and 30-minute level. Finally, the boxplot in panel (j) suggests a slight improvement with respect to the previous scenarios: the correlation coefficient for the first 10 lags is negative and increasing for the majority of our Monte-Carlo simulations. Nonetheless, since the ‘whiskers’ of the plot (denoting the $\pm 1.5 \cdot IQR$ markers of the underlying distribution) are very spread apart, we conservatively consider SF4 as not matched.

The rightmost column of Table 4 suggests that many of our objective the stylised facts are successfully reproduced by this version of the model.

4.4 Sensitivity analysis

In this section we briefly discuss the effect of varying the main parameters of the model. When not specified otherwise we intend that the effect applies across the three scenarios analysed earlier.

σ_F^2 : In the limit $\sigma_F^2 \rightarrow 0$ the price series exhibit a unit root (cf. scenario NT). Intuitively, when fundamentalists’ weights are small the prices are less ‘anchored’ to the fundamental value. On average, the larger σ_F^2 , the greater the average extent of fundamentalist behaviour, and the more often prices cross the fundamental value.

σ_C^2 : In the limit $\sigma_C^2 \rightarrow 0$, assuming $\mu_C = 0$, the persistence of returns is completely random, depending only on stochastic process ε_t . When $\sigma_C^2 > 0$ persistence is driven by both trend-followers and contrarians, giving rise to volatility clustering. Given fundamentalist parameter σ_F^2 , there exist a threshold of σ_C^2 below which the dynamics is stable, i.e. it is bounded within a neighbourhood of the fundamental price, and above which the dynamics is unstable and diverges either towards zero or $+\infty$.

μ_C : When $\mu_C > 0$ trend-followers are ‘stronger’ than contrarians. This either introduces or bolsters persistence in inter-trade times, order-flow, volatility, and volumes.

Moreover, when in conjunction with the endogenous activation scheme, it increases the number of trades because the price dynamics is more volatile and consequently more traders are activated, on average, given δ_t . Larger values of μ_C may either cause the price dynamics to become unstable (the reasoning goes as for parameter σ_C^2 above) or (raw) returns to be serially correlated, which we want to avoid. Conversely, negative values $\mu_C < 0$ wipe out all persistence in both intra-trade times and order-flow.

δ_t and ϕ : These two parameters regulate in conjunction both the total amount of trade that takes place in the market, for which they act as substitutes of one another, and the timing structure thereof. In the limit $\delta_t \rightarrow \infty$ activation is never endogenous and the average amount of trade in the market, *ceteris paribus*, is a monotone increasing function of ϕ . If $\phi = 0$ then trivially no trader is ever activated and no trade takes place; if instead $\phi = 1$ the amount of trade is maximised: exactly one trader is activated in every period (cf. scenarios NT and FC) and both volume and count of transactions are bounded from above by t . Since the (exogenous) baseline activation scheme is i.i.d., regardless of the value of ϕ , when $\delta_t \rightarrow \infty$ inter-trade times exhibit no serial correlation. In the opposite limit $\delta_t = 0$ all traders are active at every time step (regardless of ϕ); trade is maximised and bounded from above by $N \cdot t$. Provided $\delta_t > 0$, lower values of δ_t cause larger crowds of traders to participate in response to a given signal; as a result, both the order-flow and trading times tend to cluster.

tick : Large tick values cause many of the realised returns to be identically zero since multiple submitted orders are likely to accumulate at the same price. Conversely, a more fine grained tick implies dispersion of the orders at different price and therefore it triggers a non-zero price change at virtually every transaction. We find however that this effect wears out very fast with aggregation and has negligible influence on our measures based on 1-minute pooled series.

h : The shorter the order validity h , the fewer the orders stored on the book at all times and ultimately the lower the number of realised trades per day. Since the order validity is also connected to the memory span of chartists (in FC and EA scenarios), the higher h , the more likely the expected return between time t and $t + h$ is large in absolute terms (although it needn't be the case, the intuition is that cumulative returns over longer periods tend to be larger than returns over short periods), the farther the limit price of newly submitted order with respect to the current price, and consequently the steeper the dynamics of the latter. In other words, the larger h , the more likely the occurrence of sharp temporary appreciations and depreciations, that namely flash booms and flash crashes. If h is less than a certain threshold, which we find to be around 100 seconds, chartists have too short memory to introduce clustering both of absolute returns and of inter-trade times; in this case none of our scenarios is able to reproduce our target stylised facts. Conversely, for high enough values of h the dynamics may turn out unstable. This is especially true within the EA scenario when h is in excess of 13000 seconds (around 3½ hours). In between the two thresholds, we

find that the statistics relative to our target stylised facts are hardly affected by the specific value of h .

5 Concluding remarks

The distinctive statistical properties that shape financial market dynamics at high frequencies have been typically attributed to the specific patterns of information release and its diffusion among the population of traders. We show that many such properties are also obtainable in a fully endogenous setting, in which fundamental news is absent and information, originating from within the financial market as the by-product of trading activity, is common knowledge. We propose a parsimonious agent-based model in which trading emerges as the consequence of differing stable beliefs on behalf of a population of heterogeneous traders. The underlying disagreement traces to the usual specification of fundamentalist vs. chartist behaviour. A novel element that we introduce is the definition of simulation time in terms of a strict schedule that we borrow from the microstructural specification of a real stock market, namely the EURONEXT. We believe this plausibly relates each tick of our numerical simulations to proper calendar time, and enables us to investigate which properties apply within a specific time-window and how they evolve at different time-scales. We also devise a simple endogenous activation scheme that encourages traders participation in an increasing fashion with realised profit opportunities.

We find that our assumptions regarding the underlying microstructure introduce a slight dependence in the series of returns, although it quickly fades away within a couple of minutes. We also find that the fundamentalist vs chartist framework is suitable for replicating the empirically validated dependence properties of returns (leptokurtosis, absence of linear autocorrelation, and volatility clustering). Nonetheless, the introduction of our endogenous participation scheme proves crucial for the emergence of the persistence character in the timing structure of market activity. Under this scenario we are able to simultaneously reproduce, along with the stylised facts just mentioned, the fat-tailed and serially correlated nature of durations between trades, and the clustering of both volumes and order-flow.

We believe that our framework could be fruitfully extended in a number of directions. By construction, our model cannot reproduce the U-shaped pattern of intra-day market activity. More stringent assumptions regarding the traders' budget constraint or the introduction of a time feedback that puts pressure on the traders in the vicinity of market closing (e.g. due to margin requirements) could reveal useful thereof. Similarly, a more structured specification of chartists' behaviour might unveil a more asymmetric response of volatility with respect to price drops and surges (leverage effect). Finally, a rigorous calibration of the parameters of the model to better match actual financial data might allow for insightful policy experiments.

References

- Aït-Sahalia, Yacine, Jianqing Fan and Yingying Li (2013). 'The leverage effect puzzle: Disentangling sources of bias at high frequency'. *Journal of Financial Economics* 109(1), pp. 224–249. DOI: [10.1016/j.jfineco.2013.02.018](https://doi.org/10.1016/j.jfineco.2013.02.018).
- Aldridge, Irene (2013). *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Wiley. DOI: [10.1002/9781119203883](https://doi.org/10.1002/9781119203883).
- Allen, Helen and Mark P. Taylor (1990). 'Charts, Noise and Fundamentals in the London Foreign Exchange Market'. *Economic Journal* 100(400), pp. 49–59. DOI: [10.2307/2234183](https://doi.org/10.2307/2234183).
- Andersen, Torben G. and Tim Bollerslev (1997). 'Intraday periodicity and volatility persistence in financial markets'. *Journal of Empirical Finance* 4(2-3), pp. 115–158. DOI: [10.1016/S0927-5398\(97\)00004-2](https://doi.org/10.1016/S0927-5398(97)00004-2).
- Arthur, W. Brian, John H. Hilland, Blake LeBaron, Richard Palmer and Paul Tayler (1997). 'Asset Pricing Under Endogenous Expectations in an Artificial Stock Market'. In: *The Economy as an Evolving Complex System II*. Ed. by W. Brian Arthur, Steven N. Durlauf and David A. Lane. Addison-Wesley. DOI: [10.2139/ssrn.2252](https://doi.org/10.2139/ssrn.2252).
- Biais, Bruno, Pierre Hillion and Chester Spatt (1995). 'An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse'. *Journal of Finance* 50(5), pp. 1655–1689. DOI: [10.1111/j.1540-6261.1995.tb05192.x](https://doi.org/10.1111/j.1540-6261.1995.tb05192.x).
- Black, Fischer (1976). 'Studies of Stock Price Volatility Changes'. *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section*, pp. 177–181.
- Black, Fischer (1986). 'Noise'. *Journal of Finance* 41(3), pp. 529–543. DOI: [10.1111/j.1540-6261.1986.tb04513.x](https://doi.org/10.1111/j.1540-6261.1986.tb04513.x).
- Bollerslev, Tim, Julia Litvinova and George Tauchen (2006). 'Leverage and Volatility Feedback Effects in High-Frequency Data'. *Journal of Financial Econometrics* 4(3), pp. 353–384. DOI: [10.1093/jjfinec/nbj014](https://doi.org/10.1093/jjfinec/nbj014).
- Bonanno, Giovanni, Fabrizio Lillo and Rosario Mantegna (2000). 'Dynamics of the number of trades of financial securities'. *Physica A: Statistical Mechanics and its Applications* 280(1-2), pp. 136–141. DOI: [10.1016/S0378-4371\(99\)00629-9](https://doi.org/10.1016/S0378-4371(99)00629-9).
- Bouchaud, Jean-Philippe, Andrew Matacz and Mark Potters (2001). 'Leverage Effect in Financial Markets: The Retarded Volatility Model'. *Physical Review Letters* 87(22), p. 228701. DOI: [10.1103/PhysRevLett.87.228701](https://doi.org/10.1103/PhysRevLett.87.228701).
- Brock, William A. and Cars H. Hommes (1998). 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model'. *Journal of Economic Dynamics & Control* 22(8-9), pp. 1235–1274. DOI: [10.1016/S0165-1889\(98\)00011-6](https://doi.org/10.1016/S0165-1889(98)00011-6).

- Campbell, John Y., Sanford J. Grossman and Wang Jiang (1993). 'Trading Volume and Serial Correlation in Stock Returns'. *Quarterly Journal of Economics* 108(4), pp. 905–939. DOI: [10.2307/2118454](https://doi.org/10.2307/2118454).
- Chiarella, Carl and Giulia Iori (2002). 'A simulation analysis of the microstructure of double auction markets'. *Quantitative Finance* 2, pp. 346–353. DOI: [10.1088/1469-7688/2/5/303](https://doi.org/10.1088/1469-7688/2/5/303).
- Cioffi-Revilla, Claudio (2002). 'Invariance and Universality in Social Agent-Based Simulations'. *Proceedings of the National Academy of Sciences of the United States of America* 99(10), pp. 7314–7316. DOI: [10.1073/pnas.082081499](https://doi.org/10.1073/pnas.082081499).
- Cont, Rama (2001). 'Empirical properties of asset returns: stylized facts and statistical issues'. *Quantitative Finance* 1, pp. 223–236. DOI: [10.1080/713665670](https://doi.org/10.1080/713665670).
- Cont, Rama (2007). 'Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models'. In: *Long Memory in Economics*. Ed. by Gilles Teyssière and Alan P. Kirman. Springer, pp. 289–309. DOI: [10.1007/978-3-540-34625-8_10](https://doi.org/10.1007/978-3-540-34625-8_10).
- Cont, Rama (2011). 'Statistical Modeling of High-Frequency Financial Data'. *IEEE Signal Processing Magazine* 28(5), pp. 16–25. DOI: [10.1109/MSP.2011.941548](https://doi.org/10.1109/MSP.2011.941548).
- Cutler, David M., James M. Poterba and Lawrence H. Summers (1989). 'What moves stock prices?' *Journal of Portfolio Management* 15(3), pp. 4–12. DOI: [10.3905/jpm.1989.409212](https://doi.org/10.3905/jpm.1989.409212).
- Engle, Robert F. (2000). 'The Econometrics of Ultra-High-Frequency Data'. *Econometrica* 68(1), pp. 1–22. DOI: [10.1111/1468-0262.00091](https://doi.org/10.1111/1468-0262.00091).
- Euronext (2017). *Trading manual for the Universal Trading Platform*. URL: https://www.euronext.com/sites/www.euronext.com/files/trading_manual_en_03_06_2017_bob_update_static_threshold_modification.pdf.
- Fama, Eugene F. (1965). 'The Behavior of Stock-Market Prices'. *Journal of Business* 38(1), pp. 34–105. DOI: [10.1086/294743](https://doi.org/10.1086/294743).
- Foster, Andrew J. (1995). 'Volume-volatility relationships for crude oil futures markets'. *Journal of Futures Markets* 15(8), pp. 929–951. DOI: [10.1002/fut.3990150805](https://doi.org/10.1002/fut.3990150805).
- Frankel, Jeffrey A. and Kenneth A. Froot (1990). 'Chartists, Fundamentalists, and Trading in the Foreign Exchange Market'. *American Economic Review* 80(2), pp. 181–185.
- Gallant, A. Ronald, Peter E. Rossi and George Tauchen (1992). 'Stock Prices and Volume'. *Review of Financial Studies* 5(2), pp. 199–242. DOI: [10.1093/rfs/5.2.199](https://doi.org/10.1093/rfs/5.2.199).
- Ghoulmie, François, Rama Cont and Jean-Pierre Nadal (2005). 'Heterogeneity and feedback in an agent-based market model'. *Journal of Physics: Condensed Matter* 17(14), S1259–S1268. DOI: [10.1088/0953-8984/17/14/015](https://doi.org/10.1088/0953-8984/17/14/015).
- Hommes, Cars H. (2006). 'Heterogeneous Agent Models in Economics and Finance'. In: *Handbook of Computational Economics, volume 2*. Ed. by Leigh Tesfatsion and Kenneth L. Judd. North-Holland. DOI: [10.1016/S1574-0021\(05\)02023-X](https://doi.org/10.1016/S1574-0021(05)02023-X).

- Jacob Leal, Sandrine, Mauro Napoletano, Andrea Roventini and Giorgio Fagiolo (2016). 'Rock around the clock: An agent-based model of low- and high-frequency trading'. *Journal of Evolutionary Economics* 26(1), pp. 49–76. DOI: [10.1007/s00191-015-0418-4](https://doi.org/10.1007/s00191-015-0418-4).
- Jain, Prem C. and Gun Ho Joh (1988). 'The Dependence between Hourly Prices and Trading Volume'. *Journal of Financial and Quantitative Analysis* 23(3), pp. 269–283. DOI: [10.2307/2331067](https://doi.org/10.2307/2331067).
- Kirman, Alan P. and Gilles Teyssière (2002). 'Microeconomic Models for Long Memory in the Volatility of Financial Time Series'. *Studies in Nonlinear Dynamics & Econometrics* 5(4), pp. 281–302. DOI: doi.org/10.2202/1558-3708.1083.
- Kluger, Brian D. and Mark E. McBride (2011). 'Intraday trading patterns in an intelligent autonomous agent-based stock market'. *Journal of Economic Behavior & Organization* 79(3), pp. 226–245. DOI: [10.1016/j.jebo.2011.01.032](https://doi.org/10.1016/j.jebo.2011.01.032).
- Kon, Stanley J. (1984). 'Models of Stock Returns—A Comparison'. *Journal of Finance* 39(1), pp. 147–165. DOI: [10.1111/j.1540-6261.1984.tb03865.x](https://doi.org/10.1111/j.1540-6261.1984.tb03865.x).
- LeBaron, Blake (2006). 'Agent-based Computational Finance'. In: *Handbook of Computational Economics, volume 2*. Ed. by Leigh Tesfatsion and Kenneth L. Judd. North-Holland. DOI: [10.1016/S1574-0021\(05\)02024-1](https://doi.org/10.1016/S1574-0021(05)02024-1).
- Levy, Moshe, Haim Levy and Sorin Solomon (1994). 'A microscopic model of the stock market: Cycles, booms, and crashes'. *Economics Letters* 45(1), pp. 103–111. DOI: [10.1016/0165-1765\(94\)90065-5](https://doi.org/10.1016/0165-1765(94)90065-5).
- Lockwood, Larry J. and Scott C. Linn (1990). 'An Examination of Stock Market Return Volatility During Overnight and Intraday'. *Journal of Finance* 45(2), pp. 591–601. DOI: [10.1111/j.1540-6261.1990.tb03705.x](https://doi.org/10.1111/j.1540-6261.1990.tb03705.x).
- Lux, Thomas (1995). 'Herd Behaviour, Bubbles and Crashes'. *Economic Journal* 105(431), pp. 881–896. DOI: [10.2307/2235156](https://doi.org/10.2307/2235156).
- Lux, Thomas (1998). 'The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions'. *Journal of Economic Behavior & Organization* 33(2), pp. 143–165. DOI: [10.1016/S0167-2681\(97\)00088-7](https://doi.org/10.1016/S0167-2681(97)00088-7).
- Lux, Thomas and Michele Marchesi (1999). 'Scaling and criticality in a stochastic multi-agent model of a financial market'. *Nature* 397, pp. 498–500. DOI: [10.1038/17290](https://doi.org/10.1038/17290).
- Lux, Thomas and Michele Marchesi (2000). 'Volatility clustering in financial markets: a microsimulation of interacting agents'. *International Journal of Theoretical and Applied Finance* 3(4), pp. 675–702. DOI: [10.1142/S0219024900000826](https://doi.org/10.1142/S0219024900000826).
- Mandelbrot, Benoit B. (1963). 'The Variation of Certain Speculative Prices'. *Journal of Business* 36(4), pp. 394–419. DOI: [10.1086/294632](https://doi.org/10.1086/294632).

- Mandelbrot, Benoit B. (1971). 'When Can Price be Arbitraged Efficiently? A Limit to the Validity of the Random Walk and Martingale Models'. *Review of Economics and Statistics* 53(3), pp. 225–236. DOI: [10.2307/1937966](https://doi.org/10.2307/1937966).
- McInish, T. H. and R. A. Wood (1991). 'Hourly returns, volume, trade size, and number of trades'. *Journal of Financial Research* 14(4), pp. 303–315. DOI: [10.1111/j.1475-6803.1991.tb00668.x](https://doi.org/10.1111/j.1475-6803.1991.tb00668.x).
- Menkhoff, Lukas (2010). 'The use of technical analysis by fund managers: International evidence'. *Journal of Banking & Finance* 34(11), pp. 2573–2586. DOI: [10.1016/j.jbankfin.2010.04.014](https://doi.org/10.1016/j.jbankfin.2010.04.014).
- Pellizzari, Paolo and Frank Westerhoff (2009). 'Some effects of transaction taxes under different microstructures'. *Journal of Economic Behavior & Organization* 72(3), pp. 850–863. DOI: [10.1016/j.jebo.2009.08.010](https://doi.org/10.1016/j.jebo.2009.08.010).
- Raberto, Marco, Enrico Scalas and Francesco Mainardi (2002). 'Waiting-times and returns in high-frequency financial data: an empirical study'. *Physica A: Statistical Mechanics and its Applications* 314(1-4), pp. 749–755. DOI: [10.1016/S0378-4371\(02\)01048-8](https://doi.org/10.1016/S0378-4371(02)01048-8).
- Shiller, Robert J. (1981). 'Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?' *American Economic Review* 71(3), pp. 421–436. DOI: [10.3386/w0456](https://doi.org/10.3386/w0456).
- Tauchen, George E. and Mark Pitts (1983). 'The Price Variability-Volume Relationship on Speculative Markets'. *Econometrica* 51(2), pp. 485–505. DOI: [10.2307/1912002](https://doi.org/10.2307/1912002).
- Taylor, Mark P. and Helen Allen (1992). 'The use of technical analysis in the foreign exchange market'. *Journal of International Money and Finance* 11(3), pp. 304–314. DOI: [10.1016/0261-5606\(92\)90048-3](https://doi.org/10.1016/0261-5606(92)90048-3).