



An agent-based model of intra-day financialmarkets dynamics

Jacopo Staccioli, Mauro Napoletano

► To cite this version:

Jacopo Staccioli, Mauro Napoletano. An agent-based model of intra-day financialmarkets dynamics. 2018. hal-03471566

HAL Id: hal-03471566

<https://sciencespo.hal.science/hal-03471566>

Preprint submitted on 8 Dec 2021

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.

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

Jacopo Staccoli

Mauro Napoletano

SCIENCES PO OFCE WORKING PAPER n° 34

EDITORIAL BOARD

Chair: Xavier Ragot (Sciences Po, OFCE)

Members: Jérôme Creel (Sciences Po, OFCE), **Eric Heyer** (Sciences Po, OFCE), **Lionel Nesta** (Université Nice Sophia Antipolis), **Xavier Timbeau** (Sciences Po, OFCE)

CONTACT US

OFCE
10 place de Catalogne | 75014 Paris | France
Tél. +33 1 44 18 54 87

www.ofce.fr

WORKING PAPER CITATION

This Working Paper:

Jacopo Staccioli, Mauro Napoletano,
An agent-based model of intra-day financial markets dynamics
Sciences Po OFCE Working Paper, n° 34

Downloaded from URL : www.ofce.sciences-po.fr/pdf/dtravail/OFCEWPWP2018-34.pdf
DOI - ISSN

ABOUT THE AUTHOR

Jacopo Staccioli Scuola Superiore Sant'Anna, Pisa, Italy
Corresponding author, **Email Address:** j.staccioli@santannapisa.it

Mauro Napoletano OFCE, Sciences Po, Paris, France
Email Address: mauro.napoletano@sciencespo.fr

ABSTRACT

We build an agent-based model of a financial market that is able to jointly reproduce many of the stylized facts at different time-scales. 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 show that such a rule is crucial to match stylized facts.

KEY WORDS

Intra-day financial dynamics, Stylized facts, Agent-based artificial stock markets, Market microstructure, High-Frequency Trading.

JEL

C63, E12, E22, E32, O4

1 Introduction

Huge improvements in information and communication technologies have substantially reduced the latency required to operate on financial markets in the last decades. This has fostered market activity at increasingly higher frequencies. Differently from traditional money managers, who generally hold their portfolio positions for a long period, ranging from a few days to even months, high-frequency traders aim at reaping profits from a large multitude of buy and sell operations that they execute *within* each trading day, rarely holding their positions overnight. These (very) short-term trading strategies have proved remarkably profitable even during periods of nearly unprecedented financial turmoil (see e.g. Aldridge, 2013). At the same time, the overall impact of these trading strategies on market dynamics is still unclear.¹ In addition, the increasing volumes of high-frequency traders in financial markets certainly impacts many of the stylized facts of intra-daily financial market dynamics. These statistical properties are still begging for a sound theoretical framework (see Cont, 2011 for a discussion).

We propose a parsimonious agent-based model of a financial market that is able to *jointly* reproduce many of the empirically validated stylized 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 (some of the milestones include Arthur et al., 1997; Levy et al., 1994; Lux, 1995, 1998; Lux and Marchesi, 2000) has proved invaluable for investigating and replicating the statistical properties of financial markets that are hardly reconcilable with the representative agent paradigm.² However, the vast majority of the proposed models typically focuses only on a subset of the whole ensemble of recognised stylized facts, and in particular on the facts 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 stylized 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 the work of Kluger and McBride (2011), who propose a model that replicates the intra-day U-shaped seasonality in market activity, i.e. the tendency 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

¹The 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. Nevertheless, a full and agreed understanding of the functioning, potential benefits, and disadvantages of high-frequency trading has yet to be reached (see also Aldridge, 2013; Jacob Leal et al., 2016; SEC, 2014).

²Some 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. See Hommes (2006) and LeBaron (2006) for a discussion. Our work aims at contributing to the latter strand of literature.

*We wish to thank Giovanni Dosi and participants to ISCEF 2018 (Paris, France), WEHIA 2018 (Tokyo, Japan), and EAEPE 2017 (Budapest, Hungary) for helpful comments and insightful suggestions. The authors acknowledge the financial support of European Union’s Horizon 2020 grant No. 640772 – Project Dolphins. The usual disclaimer applies.

emergence of all the stylised facts that shape financial dynamics, i.e. including those at the intra-day level. We therefore attempt at filling this gap, by proposing at the same time a methodological solution to the time mapping problem and by identifying the building blocks of the model which are responsible for the emergence of the solicited stylized facts.

Our model relies on three main ingredients. The first consists of a behavioural specification of traders, which is typical of many established models in the literature. Indeed, 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 across time and common knowledge), by buying the asset if it is undervalued and selling it if it is overvalued. Chartists instead rely on the recent history of price changes to set their orders, by 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 realistic scheduling of trading events. More precisely, 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. Imposing a strict and realistic schedule on the unfolding of events enables to devise a sound and plausible correspondence between simulation iterations (which we shall identify with seconds) and calendar time. Micro-structure details about the central order book also comply 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 whenever the price change (of either sign) realised in the immediate past is high enough. The intuition 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), Aloud et al., 2013 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.³

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.

³Another 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 Teyssi re, 2002; Lux and Marchesi, 2000, and for a discussion Cont, 2007). Complementary simulation analyses that we carried out however indicate 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.

2 Stylised facts

We begin by describing the various statistical properties that characterise the intra-day dynamics of many financial markets and that our model aims at reproducing. Some of these properties are recognised to apply across different time scales while others are intra-day specific and thus require a proper calendar setting to be analysed (see Cont, 2001, 2011, for a more detailed account). The former properties are mainly related to asset returns and have already been studied and replicated in a number of agent-based models lacking of a rigorous definition of calendar time. The latter instead require a more explicit architecture in terms of microstructure. In what follows, we present each stylized fact, distinguishing between time-invariant and intra-daily facts.

2.1 Time-invariant stylized facts of financial markets

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 for derivatives pricing (see e.g. Hull, 2017).

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 – 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.

SF6 – 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.

2.2 Intra-daily stylized facts of financial markets

SF7 – Number of price changes per day In a cross-section perspective, the number of price-changing 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; McInish and Wood, 1991).

SF8 – 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 (Cont, 2011).

SF9 – 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).

SF10 – 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).

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 exception of the intra-day volume seasonality⁴ (SF11), which is unobtainable by construction in our setting as will be clear later, and of 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, the dynamics of prices and returns is endogenously determined by the interaction of the two strategies with the market microstructure, and observed 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 (Euronext, 2017):

- 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';

⁴Kluger and McBride (2011) provide an agent-based model that reproduces the U-shaped nature of intra-day 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 in the last phase⁵, 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 (see the next [section](#)). They proceed in forming their expectations about the 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.

3.2 Traders’ participation

We devise two alternative mechanisms for traders participation, one exogenous and one endogenous. In the first, a single randomly selected trader is activated at each time step. 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 exogenous mechanism is useful as a baseline scenario to describe and test the functioning of the model, we discover that the second endogenous mechanism is better suited for replicating our stylized facts. In particular, the endogenous activation allows for both crowded and uneventful periods in which either many or no orders are submitted, and contribute to clustering of volumes, of trade durations, and of the order-flow.⁶

⁵We 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.

⁶Pellizzari and Westerhoff (2009) introduce a similar rule, based on past profits.

3.3 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. The variable $p^F > 0$ denotes the fundamental price of the security, which is common knowledge among all traders. The term $h \in \mathbb{N}_+$ 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 coefficients $w_i^F \sim |\mathcal{N}(0, \sigma_F^2)|$ and $w_i^C \sim \mathcal{N}(\mu_C, \sigma_C^2)$ are trader-specific and capture the “aggressiveness” of the underlying strategy. More specifically, the weight w_i^F quantifies how quickly the price of the stock is expected to revert to its fundamental value. In contrast, the weight w_i^C measures the extent to which traders believe the future return over period h will match its past figure. From eq. (3.2), it is also clear that all chartists use only the last realised return over the time-span h to form their expectation.

The above assumptions about chartists' expectations help in containing the dimensionality of the model⁷ and stand at variance with previous works (e.g. Pellizzari and Westerhoff, 2009), which instead assume a weighted moving average (typically exponentially or linearly) over *multiple* past returns. However, given our intra-daily 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, $\ell_{i,t}$, is a tuple {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

⁷The 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.

(either fundamentalist or chartist) at time t takes the form:

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

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, which takes value 1 if the expected return is positive, -1 if it is negative, and zero otherwise.

We do not model order cancellation as an element of a trader's strategy. However, we introduce 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 the time 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 unable to reproduce 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 to come up with a closed form solution. We thus follow the standard practice in agent-based models of numerically simulating the system and then performing the relevant statistical analysis on the generated time-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, $p_0 = p^F$, 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 .

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 homogeneous time windows of one calendar minute each.⁸ 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 simulate the model⁹ under three different scenarios. In [Section 4.1](#) we only include purely noise traders. This allows us to evaluate the impact of market microstructure on the generation of stylized facts. In [Section 4.2](#) we investigate the effects of the interplay between fundamentalists and chartists on the one hand, and market microstructure on the other hand, under the baseline exogenous activation. Finally, in [Section 4.3](#) we add a further element of complexity by assuming that traders follow the endogenous participation scheme described in [Section 3.2](#). Finally, in [Section 4.4](#) we perform some complementary sensitivity analyses. 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

The first simulation scenario, which we dub NT, is useful to properly disentangle the effects implied by the market microstructure details on the generation of market statistical properties from those implied by our assumptions about traders' 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 .

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 log-differences. The average number of price changes per day under this scenario is 13578.

⁸This is necessary because, 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 two crossing orders are stored on the central order book.

⁹The simulation 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.

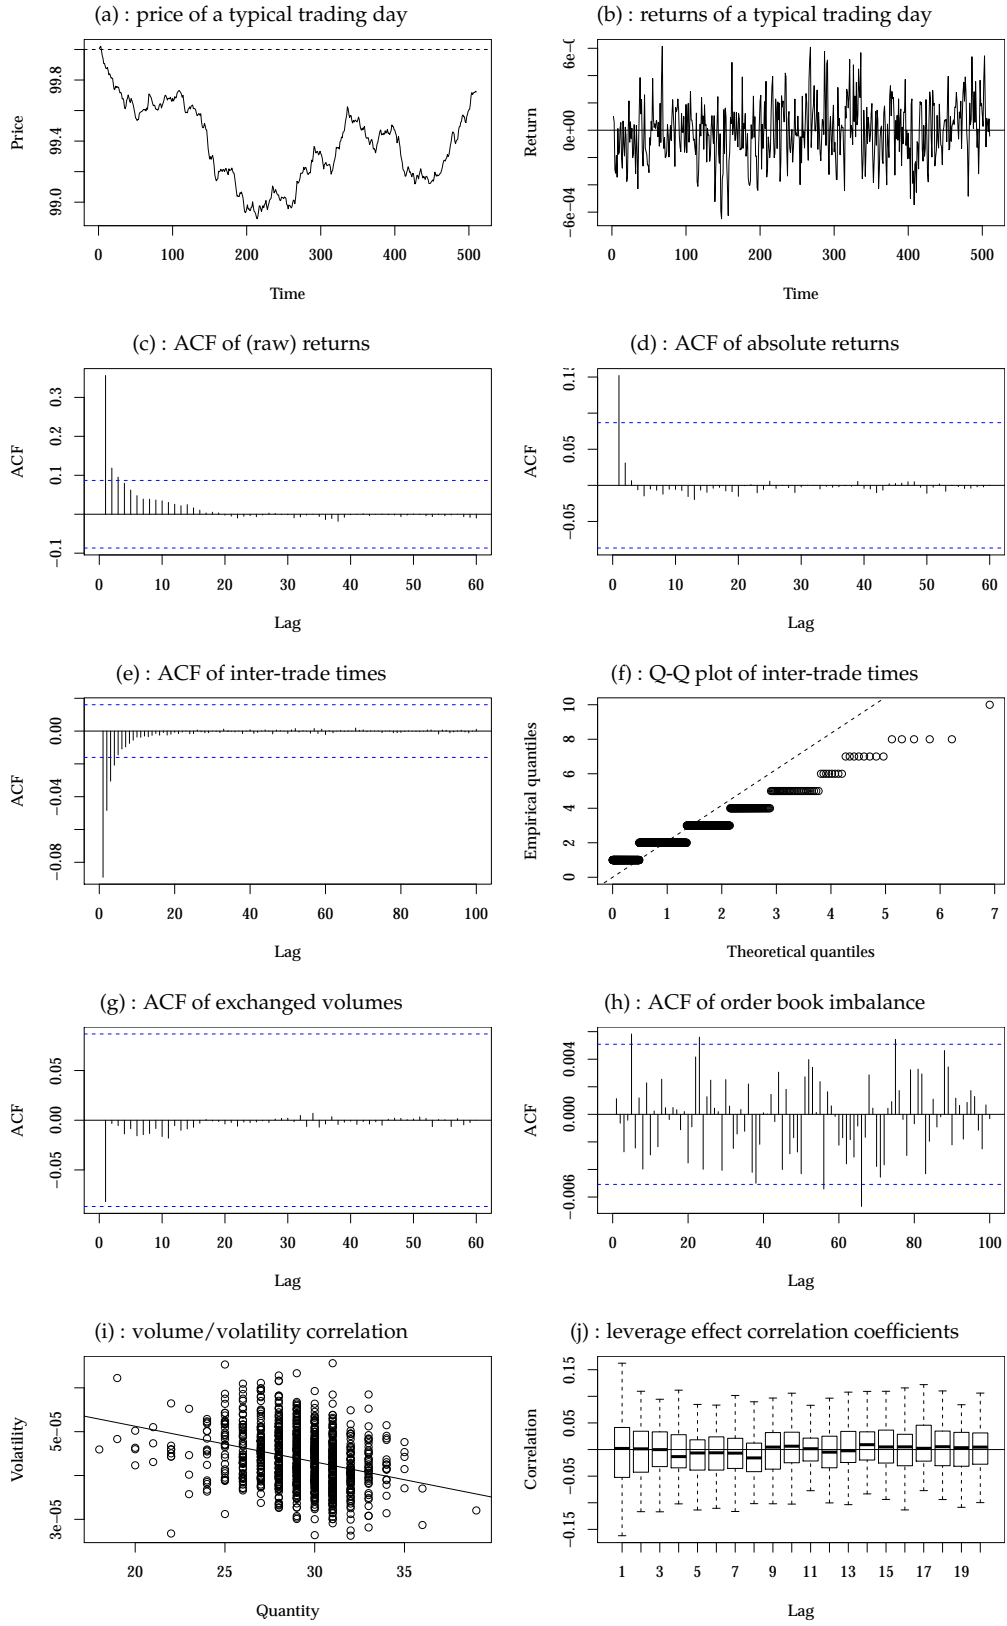


Figure 1: Main stylised facts under the noise traders scenario.

Parameter	Value
number of traders	$N = 1,000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1,000$
noise process	$\varepsilon_t \sim \mathcal{N}(\mu_\varepsilon = 0, \sigma_\varepsilon^2 = 5e-5)$
fundamentalist weight	$w_i^F = 0$
chartist weight	$w_i^C = 0$
activation threshold	$\delta_t \rightarrow +\infty$
activation fallback probability	$\phi = 1$

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

If microstructure effects were completely irrelevant, given our limit price function (3.3), then the time series of realised returns should share the same statistical properties of the i.i.d. series of expected returns. In contrast, 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 price returns. 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 (absolute) kurtosis of the sample distribution of returns, $\kappa \approx 3.43$, is only negligibly higher than that of expected returns, that by construction equals 3. We conclude that the EURONEXT microstructure setup does force a time dependence character into the resulting series, although this lasts for just under a couple of minutes.

Furthermore, 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 in this scenario. Panels (e) and (f) relate to the properties of time durations between subsequent trades. Panel (e) shows the autocorrelation function of such durations, whereas 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. This suggests 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)). Finally, panel (i) shows the presence of a negative relationship between exchanged volumes and volatility, instead of the predicted positive correlation. Likewise, panel (j) suggests that any leverage effect is absent in our series.

The first column on the right of Table 4 compares these results with our objective stylised facts. Only two facts are matched as a consequence of the interactions implied by the market microstructure. It is evident that more structure on the behaviour of the traders is needed in order to obtain a more realistic dynamics.

Parameter	Value
number of traders	$N = 1,000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1,000$
noise process	$\varepsilon_t \sim \mathcal{N}(\mu_\varepsilon = 0, \sigma_\varepsilon^2 = 5e-5)$
fundamentalist weight	$w_i^F \sim \mathcal{N}(\mu_F = 0, \sigma_F^2 = 0.001) $
chartist weight	$w_i^C \sim \mathcal{N}(\mu_C = 0.01, \sigma_C^2 = 0.1)$
activation threshold	$\delta_t \rightarrow +\infty$
activation fallback probability	$\phi = 1$

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

4.2 Fundamentalists and chartists

In this scenario (FC), we move a step forward by switching on our fundamentalist and chartist specifications, according to 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 to counteract the prevailing trend, depending on their being trend followers or contrarians. Accordingly, the stronger the magnitude of trend following behaviour, the wider the divergence of price from p^F (either upwards or downwards) should be. The parametrisation we propose, reported in Table 2, yields a price dynamics characterized by a unit root, as in the previous section (the ADF test doesn't reject the null with p -value < 0.001). Note that we set the value of $\mu_C > 0$. Reasonably, the overall sentiment among the crowd of chartists generates a self-reinforcing dynamics, rather than a self-opposing one.¹⁰

Fig. 2 shows the relevant plots under this scenario. As expected, the evolution of the price series is more “centred” around the fundamental value p^F with respect to the NT scenario thanks to the fundamentalists’ anchoring behaviour (panel (a)). However, 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).¹¹ 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 price changes, 14187, is in line

¹⁰Moreover, 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 stylised fact SF9.

¹¹In 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.

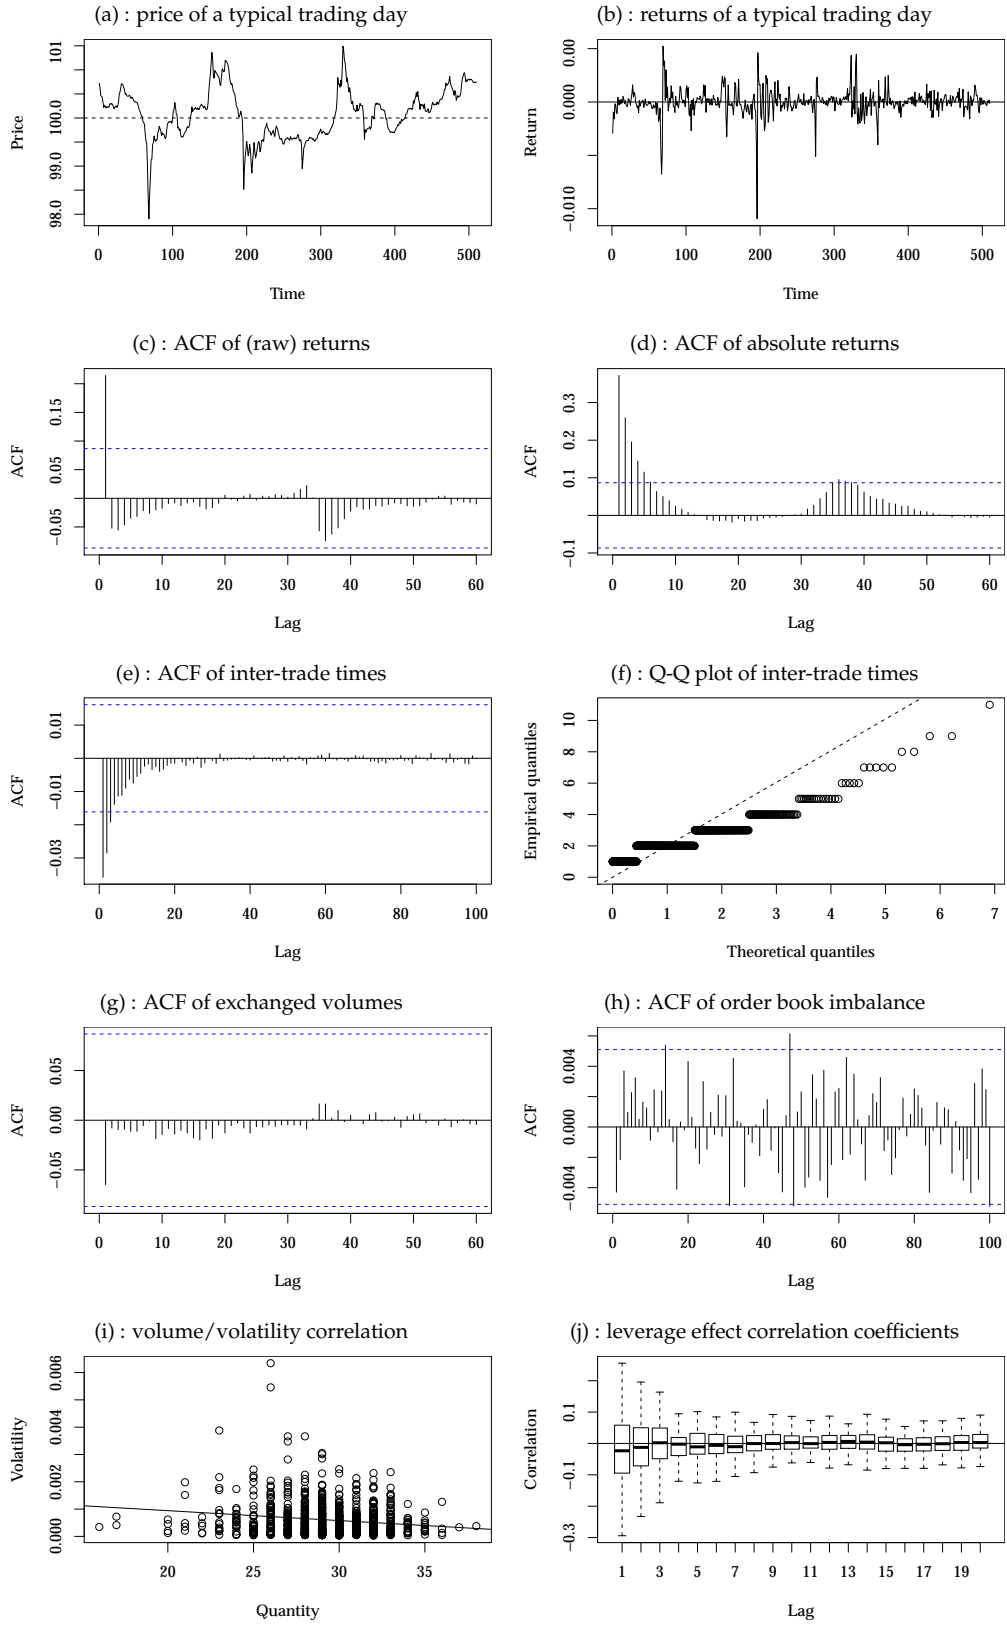


Figure 2: Main stylised facts under the fundamentalists vs. chartists scenario.

with the previous scenario. In contrast, the kurtosis of minute returns increases to 17.45, thus replicating SF1.¹²

Finally, the statical properties reported in panels (e), (f), (g), (h), (i), and (j) are qualitatively similar to the NT case. This indicates that 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 of Table 4 summarizes the list of stylized facts reproduced with the introduction of fundamentalist and chartist strategies. The improvement with respect to the noise traders scenario is clear: volatility clustering and leptokurtosis of price returns are now correctly matched. However, more structure is needed if one wants to reproduce also orders' timing and clustering properties.

4.3 Endogenous activation

In this final scenario, which we label EA, we assume that fundamentalists and chartists endogenously activate according to the scheme outlined in Section 3.2. Table 3 summarises the specific parametrisation that we employ in this scenario. The ultimate goal is to retain the properties encountered in the previous scenarios and, in addition, to replicate those properties related to the duration and clustering of orders and those about the volumes of trade.

The endogenous activation scheme captures one ever more common high-frequency nature of financial markets (see Easley et al., 2012). A crowd of traders, many of which are algorithmic machines, typically responds very quickly to a newly posted signal and engages in trading for a while until coordination on a new price has emerged.¹³

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 about price and returns generated in the latter setting have not been compromised by the new activation assumption. Leptokurtosis has increased to a minute-by-minute figure of $\kappa \approx 72.33$, decreasing to around 6.5 for 15-minute returns, and 3.7 for 30-minute returns. The average number of price changes, 11272, has decreased as a result of the new participation scheme, but is still a perfectly acceptable level for liquid traded securities (Cont, 2011).

The main benefits of endogenous activation 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 (SF8), and the quantile-quantile plot in panel (f) suggests that the tail of their distribution is fatter than exponential (SF9). Moreover, both volumes (panel (g)) and order-flow (panel (h)) are clustered (matching respectively SF5 and SF10). Panel (i) shows a positive and significant relationship between volumes and volatility (p -value

¹²The kurtosis decreases with the time window and reverts back to 3, i.e. to statistical *normality*, for 15-minute returns.

¹³In 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*.

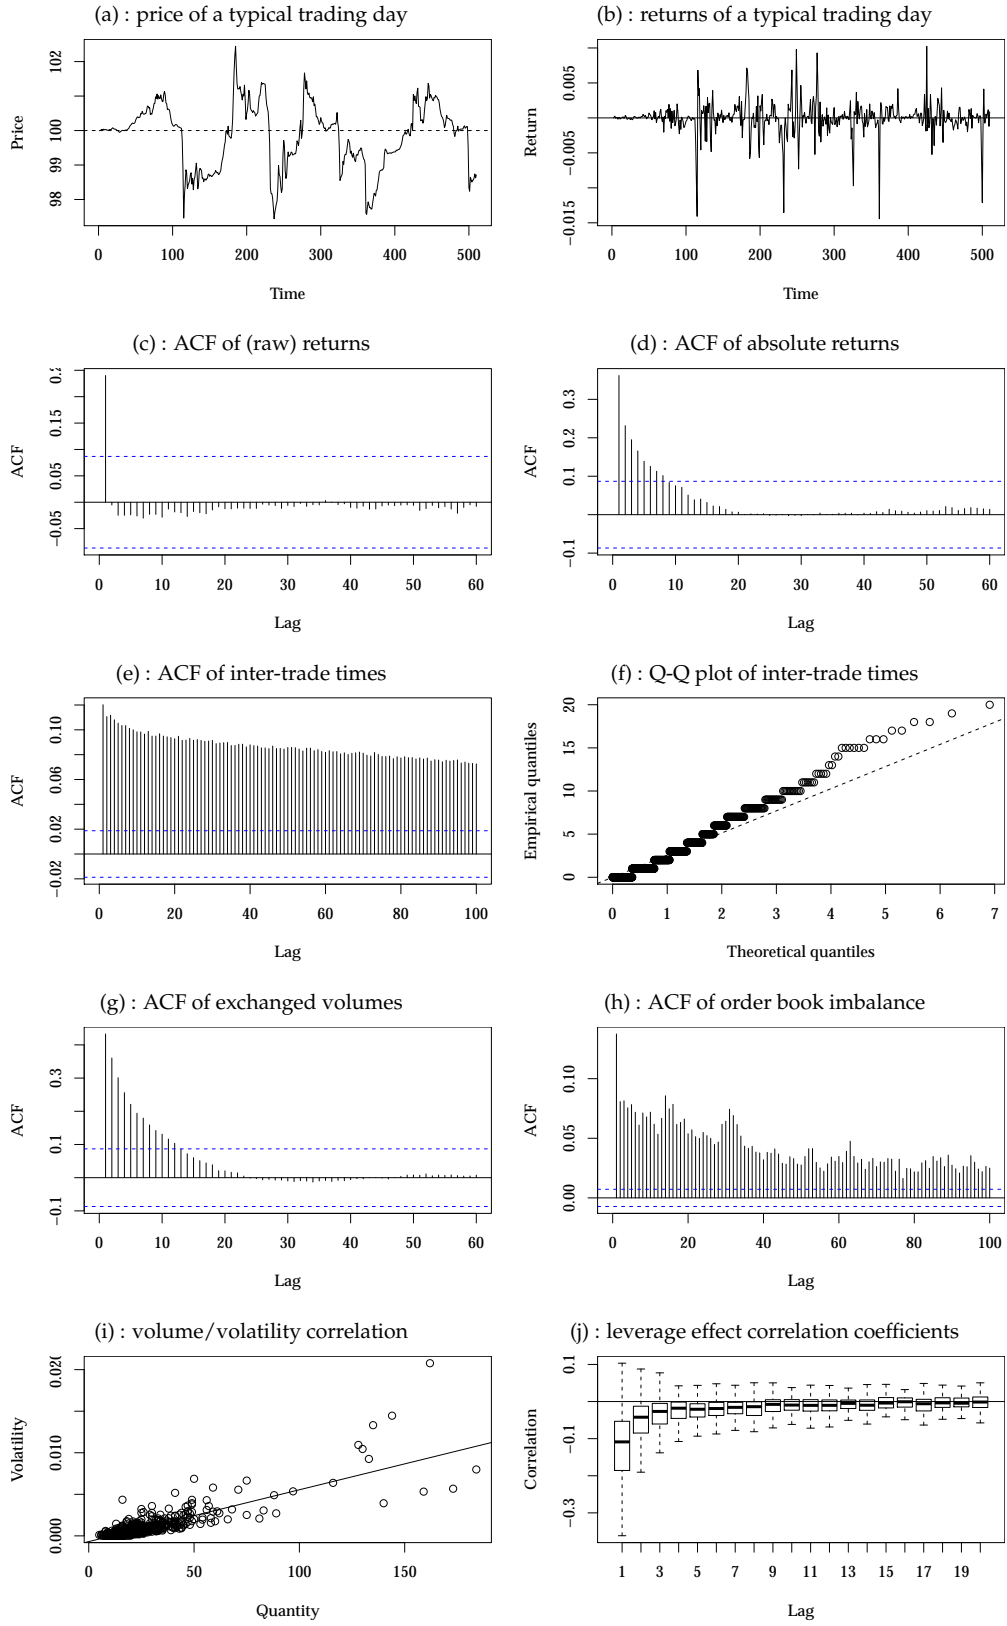


Figure 3: Main stylised facts under the fundamentalists vs. chartists scenario with endogenous activation.

Parameter	Value
number of traders	$N = 1,000$
fundamental price	$p_F = 100$
initial price	$p_0 = p_F$
smallest price change	$tick = 0.001$
horizon/order validity	$h = 1,000$
noise process	$\varepsilon_t \sim \mathcal{N}(\mu_\varepsilon = 0, \sigma_\varepsilon^2 = 5e-5)$
fundamentalist weight	$w_i^F \sim \mathcal{N}(\mu_F = 0, \sigma_F^2 = 0.001) $
chartist weight	$w_i^C \sim \mathcal{N}(\mu_C = 0.01, \sigma_C^2 = 0.1)$
activation threshold	$\delta_t \sim \mathcal{N}(\mu_\delta = 0, \sigma_\delta^2 = 0.3) $
activation fallback probability	$\phi = 1/3$

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

< 0.001) (as per SF6). 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 most of the stylized facts described in Section 2 are successfully reproduced by this version of the model. In particular, the emergence of the properties about orders’ duration and clustering is very much linked to the dynamics induced by endogenous activation in this scenario. Indeed, the level of the variance of the distribution of agents’ activations thresholds, $\sigma_\delta^2 = 0.3$, 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 price adjustment process.

4.4 Sensitivity analysis

In this section we briefly discuss the effect of varying, one at a time, the main parameters of the model¹⁴ in a neighbourhood of the parametrisation used in the most complete scenario, i.e. EA scenario, with endogenous traders’ activation (cf. Table 3). For each of the discussed

¹⁴We also experimented with changes in other parameters of the model besides the ones discussed in this section. We do not report the results of these additional sensitivity analyses here. However, they are available from the authors upon request.

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	autocorrelation of volumes	X	X	✓
SF6	volume/volatility correlation	X	X	✓
SF7	number of price changes per day	✓	✓	✓
SF8	autocorrelation of durations	X	X	✓
SF9	fat-tailed durations	X	X	✓
SF10	order-flow clustering	X	X	✓
SF11	U-shaped activity	X	X	X

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

parameters, we present plots showing the change in the relevant statistics of the simulation and we relate it to the stylized facts of financial markets.

Changes in the variance of the fundamentalists' weights, σ_F^2

In the limit of $\sigma_F^2 \rightarrow 0$, the price becomes less “anchored” to the fundamental value. In this extreme case the persistence of volatility (absolute returns), of volumes, of trade durations, and of the order flow, the number of price changes, and the leptokurtic signature of returns are maximised. The plot in panel (a) of Fig. 4 shows the number of statistically significant lags (at the 95% confidence level) as a function of σ_F^2 for each of the autocorrelation functions calculated in the EA simulation scenario (cf. Fig. 3). Autocorrelations computed on minute-by-minute data (namely returns, absolute returns, and volumes) refer to the left scale. Those computed on tick data (namely durations and order flow) are on the right scale. In addition, panel (b) depicts the effect of σ_F^2 upon the number of price changes per day (left scale) and the minute returns' kurtosis (right scale). All measures are averages across 100 Monte Carlo simulations for each value of the underlying parameter.

From the analysis of the above mentioned figures it is clear that as σ_F^2 grows, fundamentalists counteract the effect of technical traders and all the considered statistics either decrease or remain unchanged. It is also clear from the picture that σ_F^2 has not effect whatsoever on the persistence of raw returns.

Changes in the variance of the chartists' weights, σ_C^2

At very low values of σ_C^2 and with $\mu_C > 0$ the distribution of chartists' weights is significantly skewed towards trend-following strategies. This causes raw returns to be autocorrelated for several lags (see Fig. 5(a), left scale). As σ_C^2 increases the distribution of the chartists' strategies tends to be more balanced between trend-followers and contrarians. This decreases the autocorrelation of raw returns without significantly influencing the persistence of either absolute returns and of traded volumes.

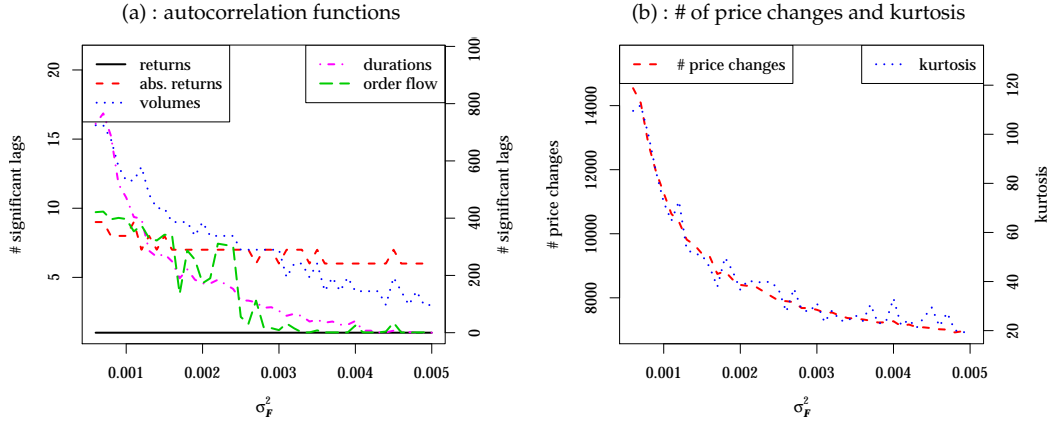


Figure 4: Sensitivity analysis for σ_F^2 . Other parameters as per the EA scenario.

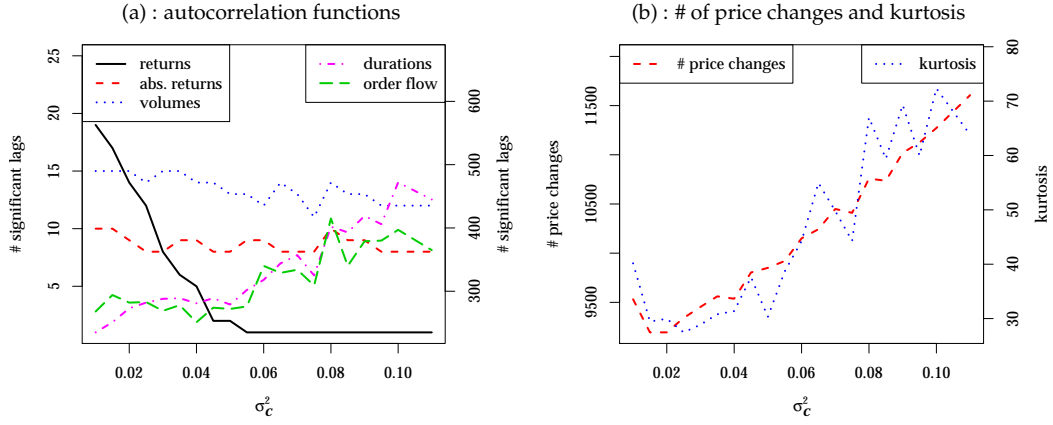


Figure 5: Sensitivity analysis for σ_C^2 . Other parameters as per the EA scenario.

At the same time, the extent of chartists' disagreement about future returns increases with σ_C^2 , because the weights of followers and contrarians are located farther and farther away from one another. *Ceteris paribus*, this increases the occurrence of abnormal returns (cf. Fig. 5(b)). In addition, abnormal returns and the mechanism of endogenous activation (see Sections 3.2 4.3) encourage several traders to post their orders simultaneously. This increases the number of price changes per day and generates persistence in inter-trade durations and in order flows (see Fig. 5(a), right scale).

Changes in the variance of traders' activation thresholds, σ_δ^2 , and of the fallback probability ϕ

These two parameters jointly regulate both the total amount of trade that takes place in the market and the timing structure thereof. First, note that, by definition $\delta_{i,t} \sim |\mathcal{N}(0, \sigma_\delta^2)|$, and therefore $\mathbb{E}[\delta_{i,t}]$ is an increasing function of σ_δ^2 . In the limit of $\sigma_\delta^2 \rightarrow \infty$, $\delta_t \rightarrow \infty$ and

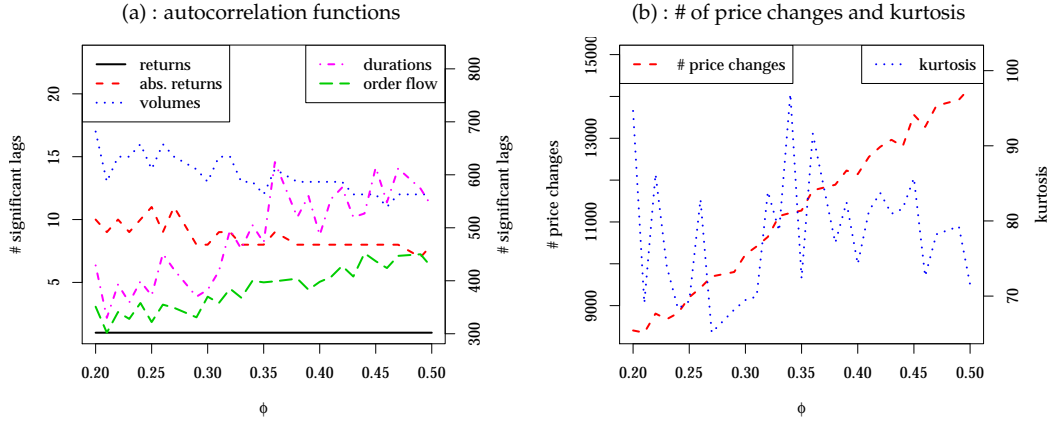


Figure 6: Sensitivity analysis for ϕ . Other parameters as per the **EA** scenario.

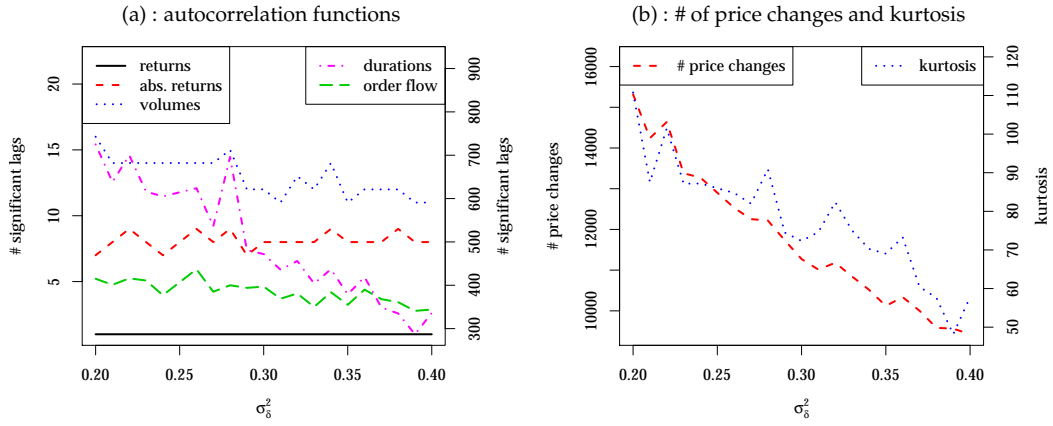


Figure 7: Sensitivity analysis for δ . Other parameters as per the **EA** scenario.

activation is never endogenous and the average amount of trade and of price changes in the market is *ceteris paribus*, a monotonically increasing function of ϕ (cf. the dashed red line in Fig. 6(b)). 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, under the assumption of uniform activation: exactly one trader is activated in every period (cf. scenarios **NT** and **FC**) and both volumes and the number of transactions (a superset of the number of price changes) are bounded from above by t .

Furthermore, when σ_δ^2 is large $\delta_{i,t}$ is also large on average. In this case inter-trade times exhibit little serial correlation (cf. Fig. 7(a)). In the opposite limit, $\delta_t = 0$, all traders are instead active at every time step (regardless of ϕ), and the number of transactions is maximised and bounded from above by $N \cdot t$. Finally, lower average values of δ_t cause larger crowds of traders to participate in response to a given signal, boosting the number of price changes (cf. Fig. 7(b)). As a result, both the order-flow and trading times tend to cluster (see again Fig. 7(a)).

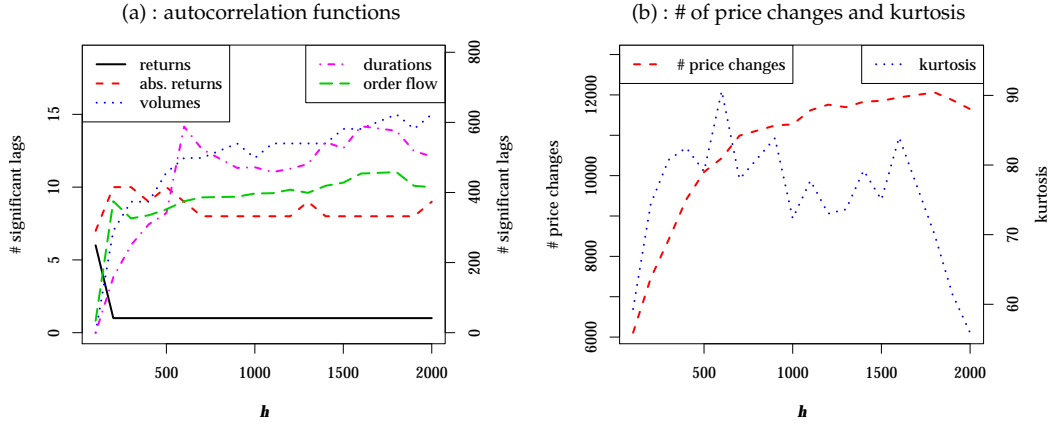


Figure 8: Sensitivity analysis for h . Other parameters as per the EA scenario.

Changes the investment horizon and orders' validity: h

The shorter the order's validity h , the fewer the orders stored on the book at all times. Accordingly, the number of price changes per day decreases as a function of h (cf. Fig. 8(b)).

Furthermore, the parameter h also sets the memory span of chartists (cf. eq. (3.2)). An increase in this parameter thus leads the latter to exacerbate small trends in prices. This is because cumulative returns over longer periods tend to be larger than returns over short periods in presence of some positive autocorrelation. Accordingly, the higher h , the more likely expected returns between time t and $t + h$ are large in absolute terms, and the farther is the limit price of newly submitted orders with respect to the current price. This effect, combined with the endogenous activation mechanism, contributes to the persistence of absolute returns, volumes, orders' duration and sides, which are all observed for small increases in h when starting from a low base value of $h = 100$ (see Fig. 8(a)). It also explains the increase in the returns kurtosis displayed in (cf. Fig. 8(b)).

Nevertheless, the above effects quickly vanish with further increases in h . Indeed, the number of significant lags in (raw) returns autocorrelation quickly decays, eventually stabilizing at a low value of 1 already for h around 100. In addition, the degrees of persistence in absolute returns and in order flow and durations also converge to a stationary value. Finally, the kurtosis of price returns evolves non-monotonically with h , eventually returning to the same low value observed at $h = 100$. Notice that in the EA scenario we set $h = 1000$. It follows that increases in the value of h in that scenario have no significant effects on the autocorrelation functions of the main market variables we consider.

5 Concluding remarks

The distinctive statistical properties that shape financial market dynamics at daily and intra-daily 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 can be simultaneously reproduced in a framework wherein fundamental news are absent and information, originating from within the financial market (as the by-product of trading activity) is common knowledge. We build a parsimonious agent-based model in which trading and its statistical properties emerge endogenously out of the interaction between fundamentalist and chartist strategies on the one hand, and a realistic market microstructure specification on the other hand.

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 iteration 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, which 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 can be fruitfully extended in several directions. First, our model cannot reproduce, by construction, 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 traders close to the end of the trading day (e.g. due to margin requirements) could be useful in this respect. 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, in this paper we only considered the ability of the model to qualitatively replicate the main statistical properties of financial markets. However, one could further fine-tune the calibration of the parameters of the model by exploiting actual financial data. This might allow one to perform quantitative experiments on regulatory policies affecting market microstructure or trading behaviour.

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: [10.1002/9781119203803](https://doi.org/10.1002/9781119203803).
- 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).
- Aloud, M., E. Tsang and R. Olsen (2013). 'Modeling the FX market traders' behavior: An agent-based approach'. In: *Simulation in Computational Finance and Economics: Tools and Emerging Applications*. Ed. by B. Alexandrova-Kabadjova, S. Martinez-Jaramillo, A. Garcia-Almanza and E. Tsang. Hershey, PA: Business Science Reference. DOI: [10.4018/978-1-4666-6268-1.ch018](https://doi.org/10.4018/978-1-4666-6268-1.ch018).
- 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).
- 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).
- Easley, David, Marcos M. López DE Prado and Maureen O'Hara (2012). 'The Volume Clock: Insights into the High-Frequency Paradigm'. *Journal of Portfolio Management* 39(1), pp. 19–29. DOI: [10.2139/ssrn.2034858](https://doi.org/10.2139/ssrn.2034858).
- 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_2018_01_03.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).
- Hull, J.H. (2017). *Options, Futures, and Other Derivatives*. 9th ed. Pearson.

- 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 Teyssi re (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 (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).
- SEC (2014). *Equity Market Structure Literature Review. Part II: High Frequency Trading*. White Paper. March 18, 2014. URL: https://www.sec.gov/marketstructure/research/hft_lit_review_march_2014.pdf.
- 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).



ABOUT OFCE

The Paris-based Observatoire français des conjonctures économiques (OFCE), or French Economic Observatory is an independent and publicly-funded centre whose activities focus on economic research, forecasting and the evaluation of public policy.

Its 1981 founding charter established it as part of the French Fondation nationale des sciences politiques (Sciences Po), and gave it the mission is to “ensure that the fruits of scientific rigour and academic independence serve the public debate about the economy”. The OFCE fulfils this mission by conducting theoretical and empirical studies, taking part in international scientific networks, and assuring a regular presence in the media through close cooperation with the French and European public authorities. The work of the OFCE covers most fields of economic analysis, from macroeconomics, growth, social welfare programmes, taxation and employment policy to sustainable development, competition, innovation and regulatory affairs.

ABOUT SCIENCES PO

Sciences Po is an institution of higher education and research in the humanities and social sciences. Its work in law, economics, history, political science and sociology is pursued through [ten research units](#) and several crosscutting programmes.

Its research community includes over [two hundred twenty members](#) and [three hundred fifty PhD candidates](#). Recognized internationally, their work covers [a wide range of topics](#) including education, democracies, urban development, globalization and public health.

One of Sciences Po's key objectives is to make a significant contribution to methodological, epistemological and theoretical advances in the humanities and social sciences. Sciences Po's mission is also to share the results of its research with the international research community, students, and more broadly, society as a whole.

PARTNERSHIP
