Journal of Economic Interaction and Coordination

Fast traders and slow price adjustments: an artificial market with strategic interaction and transaction costs.

--Manuscript Draft--

Manuscript Number:	JEIC-D-18-00009R2
Full Title:	Fast traders and slow price adjustments: an artificial market with strategic interaction and transaction costs.
Article Type:	S.I. : WEHIA 2017
Corresponding Author:	Danilo Liuzzi Universita degli Studi di Milano ITALY
Corresponding Author Secondary Information:	
Corresponding Author's Institution:	Universita degli Studi di Milano
Corresponding Author's Secondary Institution:	
First Author:	Danilo Liuzzi
First Author Secondary Information:	
Order of Authors:	Danilo Liuzzi
	Marco Tolotti
	Paolo Pellizzari
Order of Authors Secondary Information:	
Funding Information:	
Abstract:	In this paper, we propose an artificial market to model high-frequency trading where fast traders use threshold rules strategically to issue orders based on a signal reflecting the level of stochastic liquidity prevailing on the market. A market maker is in charge of adjusting prices (on a fast scale) and of setting closing prices and transaction costs on a daily basis, controlling for the volatility of returns and market activity. We first show that a baseline version of the model with no frictions is able to generate returns endowed with several stylized facts. This achievement suggests that the two time scales used in the model are one (possibly novel) way to obtain realistic market outcomes and that high-frequency trading can amplify liquidity shocks. We then explore whether transaction costs can be used to control excess volatility and improve market quality. While properly implemented taxation schemes may help in reducing volatility, care is needed to avoid excessively curbing activity in the market and intensifying the occurrence of abnormal peaks in returns.
Response to Reviewers:	Dear Editor, we submit a deeply revised version of our work and thank you for the request to polish once more the text. We decided to use a professional editor and we have indeed fixed many typos and improved the language in (literally) dozens of places. We truly believe the paper is now more readable and clearer. Do not hesitate to contact us if any further action is needed. Best regards, Danilo Liuzzi, Paolo Pellizzari and Marco Tolotti

Fast traders and slow price adjustments: an artificial market with strategic interaction and transaction costs

Danilo Liuzzi*[1], Paolo Pellizzari[2], and Marco Tolotti[1]

[1]Department of Management, Ca' Foscari University of Venice
[2]Department of Economics, Ca' Foscari University of Venice

January 30, 2018

*Corresponding author. Address: Cannaregio 873, 30121 Venice, Italy; email: *danilo.liuzzi@unive.it*

Fast traders and slow price adjustments: an artificial market with strategic interaction and transaction costs

September 26, 2018

Abstract

In this paper, we propose an artificial market to model high-frequency trading where fast traders use threshold rules strategically to issue orders based on a signal reflecting the level of stochastic liquidity prevailing on the market. A market maker is in charge of adjusting prices (on a fast scale) and of setting closing prices and transaction costs on a daily basis, controlling for the volatility of returns and market activity.

We first show that a baseline version of the model with no frictions is able to generate returns endowed with several stylized facts. This achievement suggests that the two time scales used in the model are one (possibly novel) way to obtain realistic market outcomes and that high-frequency trading can amplify liquidity shocks. We then explore whether transaction costs can be used to control excess volatility and improve market quality. While properly implemented taxation schemes may help in reducing volatility, care is needed to avoid excessively curbing activity in the market and intensifying the occurrence of abnormal peaks in returns.

Keywords: Artificial markets; High-frequency trading; Liquidity shocks; Transaction costs.

1 Introduction

Financial markets aggregate the beliefs and trading decisions of a myriad of agents endowed with different objectives, strategies, information, and abilities. The stunning complexity of the outcomes is revealed in the nontrivial properties of financial returns that feature a set of intriguing and almost universal statistical properties known as stylized facts.[1] For a review, see Lux (2009) or the evergreen Cont (2001).

[1]This terminology, now quite common in the financial literature but possibly less used elsewhere, was actually introduced by the economist Nicholas Kaldor to refer to the most relevant elements requiring explanation. In his words, "The theorist... ought to start off with a summary of the facts which he regards as relevant to his problem [and] concentrate on broad tendencies, ignoring individual detail, and proceed on the 'as if' methods, i.e. construct a hypothesis that could account for these 'stylized facts'...", see Kaldor (1961).

In the last two decades, many models have to some extent been able to generate a handful of such regularities, suggesting possible sufficient drivers for the presence of stylized facts.

In this paper, we explore a relatively new avenue and describe an artificial market where actions take place over two time scales: at the daily level, a market maker (or some entity/organization in charge of running the exchange) adjusts the price based on the excess demand for a risky asset and on some adjustable transaction cost (or tax). The closing price crucially depend also on a slowly varying market depth stochastic process, which can be thought of in terms of exogenous fluctuating (inverse) liquidity of the market. On the intra-day time scale, a large number of traders interact, crudely aiming at maximizing their short-term returns (net of transaction costs). Given the furious pace of this high-frequency trading, traders make their decisions using fast rules based on activation thresholds. These thresholds trigger sales or purchases that are contingent on a heterogenous individual signal and on an educated guess of the direction the price will take due to the decisions of all other agents.

The market maker mechanically mediates fast trades and sets trading costs based on the prevailing liquidity level with no explicit or modelled objective. In contrast, the fast agents in the model strive to ideally find a local-in-time Nash equilibrium where everyone optimally buys/sells the asset given his/her noisy individual signal and actions of other traders. The result of this mechanism design is a sequence of frequent (discrete) adjustments of the intra-day price for the risky asset. In our opinion, this modeling assumption, which resembles the idea of *frequent batch auctions* recently discussed in Budish et al. (2015), is a simple but realistic representation of the market: high-frequency traders have some sophistication, use private signals, do not want to be outsmarted, and have to decide using fast rules in a sequence of best responses to other traders who quickly approximate a (local) pricing equilibrium. The resulting end-of-day price is finally determined by a market maker who takes into account the liquidity of the market once a day and sets the appropriate trading costs (or, in an alternative interpretation, levies a transaction tax).

Our results are threefold. Firstly, the returns of our artificial market are endowed with fat tails, sizeable (excess) kurtosis, no linear predictability, and some volatility clustering. Another important result is the observation that this dual-scale mechanism, where strategic interaction happens at the intra-day level, can intensify and magnify liquidity shocks occurring over much longer (daily) time scales. Thirdly, we show that transaction costs have a non-trivial impact on the statistical properties of returns: while there is some potential to neutralize liquidity shocks, increased frictions appear to reduce volatility and activity, but generate spikes in prices and larger kurtosis in returns.

Our paper takes into account the increasing awareness that (substantial) heterogeneity is needed in models of financial markets to obtain more realistic returns, see Lux (2009) and Kirman (1992) for a critical discussion of the representative agent approach. In our setup, traders' decisions are based on heterogeneous thresholds as well as on strategic interaction with other agents, as pioneered in Granovetter (1978). Several papers in the last two decades have linked financial markets with dynamics arising from simple models where basic 'particles' influence each other at the individual level to produce interesting aggregates, often taking the limit for a large number of traders. While it may still be a challenge to provide a canonical microfoundation of such interactions, many results are of interest for financial economists and deserve wider recognition; see the survey provided in Lux (2016).

Another stream of literature has influentially pointed out how a detailed microstructural representation of exchange may play an important role in obtaining realistic returns, or may even be responsible for the presence and intensity of several stylized facts and large jumps; see, for instance, Chiarella et al. (2009) and Maslov (2000). The model presented in this paper differs from previous works in that many details of current intra-day markets are abstracted away to focus on the strategic search for a local pricing equilibrium. As such, we incorporate some form of herding effect and shared sentiment in the fast component of our model, as done in distinct ways in LeBaron and Yamamoto (2008) or Chiarella et al. (2017). In order to obtain our distributional results, however, there is no need to include a detailed implementation of a continuous double auction that would most likely reinforce our findings. Our model may also suggest that high-frequency trading can create lively price dynamics even if, as in our setup, returns are ultimately settled and smoothed by a conservative market maker on a daily basis; see Hasbrouck and Saar (2009, 2013) for insightful descriptions of how fast trading and technological innovations impact traditional market models.

The paper is organized as follows. In Section 2, we define the model and its two basic components, namely how agents are involved in high-frequency trading. We also provide a description of the market maker who, on a daily basis, sets the closing price and adjusts the prevailing transaction costs, generally accounting for an exogenous liquidity measure. The third section is devoted to an analysis of returns: we outline the procedure used to determine the parameters of the model with the aim of matching some of the most significant stylized facts known in the literature; we then discuss the relationships between return/prices and liquidity in the market; we ultimately propose two different mechanisms for introducing a non-null transaction cost and describe the resulting impact on returns and market quality. Section 4 summarizes and draws up some closing remarks.

2 An artificial market for high-frequency trading

In this section we present a stylized artificial market for high-frequency trading. Agents can buy or sell a share of a risky asset and base their action upon the forecast of short-period returns. A market maker, or another external authority, sets (daily) prices and possibly exacts a transaction cost on returns.

Since we are aiming at modeling *fast* trading and the relative emerging *daily* prices, we introduce two time scales: a *fast time scale* for intra-day activities (revise trading strategies, adjust prices) and a *slow time scale* for daily activities (compute daily returns, set daily transaction costs). We use $n \in N$ to denote the daily calendar dates: for each n, we assume that a random number τ_n of intra-day time steps take place: $t = 1, \ldots, \tau_n$. Put differently, each time the market moves (i.e. we have some activities), the fundamentals of the market change and we account for this using the intra-day (fast) time scale. Finally, we assume that $(\tau_n)_{n\in N}$ is a sequence of i.i.d. random numbers extracted from a Poissonian distribution.

In the remainder of this section, we describe how trading occurs in the market, determining the demand for the asset and the intra-day (internal) returns of each trading operation; finally, we discuss how the time series of returns and prices are influenced by changing the level of taxation on the market.

2.1 High-frequency trading on a fast intra-day time scale

Fast trading requires simple rules to determine an order for buy or sell operations. In particular, we focus on a simple threshold: if the expected internal return of a *buy operation* is large enough, say bigger than a certain threshold level T^b, the action is *Buy*; if the return is smaller than a lower threshold T^s, the operation is *Sell*. Otherwise, the agent does not activate any operation. Thresholds T^b and T^s depend on the expected return $R_n(t)$ being realized in the short run by a purchase/sale at the intra-day period t and calendar date n.[2]

Having this baseline paradigm in mind, we rely on the stylized artificial market proposed in Fontini et al. (2016) to model market activities. On the market, I agents are active and can trade one share of the asset at discrete intra-day time steps. To keep the mechanism as simple as possible, we assume that, at each trading time step, agents may own at most one share of the asset. We denote the *ownership* state variable by $\omega_i(t) \in \{0, 1\}$, for $i = 1, \dots, I$: $\omega_i(t) = 1$ means that agent i owns the share on $(t, t+1]$; if not, $\omega_i(t) = 0$ means that agent i owns no share of the asset. They bet on which return $R_n(t)$ is going to prevail in the next trading period, where the return is proportional to the excess demand for the asset:[3]

$$Q_n(t) = \frac{1}{I}\sum_{j=1}^{I}\omega_j(t), \text{ for } t = 1, \dots, \tau_n \tag{1}$$

In more detail,

$$R_n(t) = k_n\,\Delta_n(t), \tag{2}$$

where k_n measures the *market depth* at date n (it can be thought of as an inverse of the liquidity available on the market) and where

$$\Delta_n(t) = Q_n(t) - Q_n(t-1)$$

is a measure of the variation in demand for the asset at each time step.

Each agent is endowed with a private signal $\epsilon(t)$, statistically and independently drawn from a common distribution η about the realized return. As commonly assumed in the standard literature on discrete choice models (see, for instance, Brock

[2] To facilitate the comprehension, we use capital letters to denote all intra-day variables, whereas small letters are used to denote daily variables. Accordingly, intra-day variables are indexed by t, whereas calendar dates are indexed by n.

[3] We assume that the demand is always fulfilled by a market maker, who adjusts taxation in order to limit/entice activity on the market.

and Durlauf (2001)), we assume that η is a centered logistic distribution with dispersion parameter $\beta > 0$, so that

$$\mathbb{P}(\epsilon < x) = \frac{1}{1 + e^{-\beta x}}. \tag{3}$$

It can be shown that the variance of the distribution is equal to $\pi^2/3\beta^2$. Therefore, the larger β is, the less dispersed agents' opinion about returns is. Finally, agents' decisions also depend on the level of transaction cost imposed by the market maker. For simplicity, we assume that this is done by deducting a predetermined term $\mu_n \geq 0$ from the realized returns.

In sum, the decision scheme works as follows. Each trader $i = 1, \dots, I$ observes μ_n and k_n, and forms an expectation $\mathbb{E}^i[R_n(t)]$ about the prevailing return.[4] Finally, he decides whether or not to invest during intra-day period t, depending on payoffs U, as described in Table 1.

$\omega_i(t-1) \to \omega_i(t)$	$U_i(\omega_i(t); \omega_i(t-1))$
$0 \to 1$ (buy)	$\mathbb{E}^i[R_n(t)] - \mu_n + \epsilon_i(t)$
$0 \to 0$ (sleep)	0
$1 \to 1$ (keep)	$\mathbb{E}^i[R_n(t)] + \epsilon_i(t)$
$1 \to 0$ (sell)	$-\mu_n$

Table 1: Scheme of the four possible actions and related payoffs.

The payoffs depend on both previous and actual actions. They summarize, for each possible situation, the monetary gain (loss) expressed in terms of returns. For example, the payoff related to a *buy operation* is related to the benefit from ownership (the expected return plus the private signal) minus the cost μ_n to be paid to enter the market. Note that μ_n is equivalent to a transaction cost proportional to the asset price that is exchanged. Similarly, the payoff related to a *sell operation* is simply given by the exit cost $-\mu_n$.The owner will be more willing to pay the cost if he forecasts a negative return larger than μ_n in absolute value.

An agent j, who does not hold the share, decides to enter the market as soon as the perceived return, net of the cost, is higher than the stay-out option (which yields null return in the period under consideration). Eventually, this happens when $\mathbb{E}^j[R_n(t)] + \epsilon_j(t) - \mu_n > 0$. This translates into an *agent-specific threshold*

$$T^b_j(t) = -\mathbb{E}^j[R_n(t)] + \mu_n,$$

triggering a purchase when $\epsilon_J(t) > T^b_j(t)$. Conversely, agent i, who owns the asset, decides to leave the market when the perceived return is smaller than the exit cost μ_n, i.e. $\mathbb{E}^i[R_n(t)] + \epsilon_i(t) < -\mu_n$. The threshold signaling to sell is, therefore,

$$T^s_i(t) = -\mathbb{E}^i[R_n(t)] - \mu_n.$$

The expected return, the random signal and the transaction cost shape the agents' final decision, expressed in terms of endogenous thresholds, as shown in Table 2.

[4] $\mathbb{E}^i[\cdot]$ is the expectation with respect to the joint distribution of vector $\underline{\epsilon^{-i}} = (\epsilon_j)_{j \neq i}$.

Original status	(state variable)	Threshold	Action
Owner	$(\,\omega_i(t-1)=1\,)$	$\epsilon_i(t) < T^s_i(t)$	*Sell*
Non-owner	$(\,\omega_j(t-1)=0\,)$	$\epsilon_j(t) > T^b_j(t)$	*Buy*

Table 2: Threshold levels and related actions.

Two remarks must be made. Firstly, the returns depend on k_n and μ_n, which are not constant but are *slowly varying* on the daily time scale. Therefore, for the purposes of fast traders, they will be considered as constant for intra-day activities. Secondly, note that $R_n(t)$ is a function of $Q_n(t)$ and, therefore, it implicitly depends on the actual trading strategy of all agents on the market. This introduces a *strategic interaction mechanism*: the single decision $\omega_i(t)$ of agent i involves his forecast of the entire vector $\underline{\omega}(t) = (\omega_1(t), \ldots, \omega_I(t))$. As a consequence, the expected returns and thresholds $\underline{T}^b$ and $\underline{T}^s$ themselves are implicitly determined, and are related to the existence of a Nash equilibrium for the game played by I agents at each trading time step t. We postpone the discussion of the existence of such thresholds and the mechanism leading to their identification to Section 2.2.

Having defined intra-day returns, it is possible to determine intra-day price values. Given an initial price $P_n(0)$ for the asset, intra-day prices are

$$P_n(t) = P_n(t-1) \cdot e^{R_n(t)}, \quad t = 1 \ldots, \tau_n. \tag{4}$$

At each date n, the market maker computes the closing price $p_n = P_n(\tau_n)$ and makes it public.[5] Similarly, daily returns are computed as $r_n = \log(p_n/p_{n-1})$, for $n \geq 1$. Therefore, daily prices corresponding to closing values and returns are determined once an initial price $p_0 > 0$ has been provided. Note that p_n (respectively, r_n) refers to daily prices (returns) and differs from P_n (R_n), which reflects the intra-day price (return).

In order to complete the description of the market mechanism, we must set k_n and μ_n. Concerning the latter, for the moment we simply assume $\mu_n \equiv 0$. We postpone a discussion about the market maker policy and its implications to Section 3.3. Concerning the former, we assume that $(k_n)_{n\in N}$ evolves as a discretized version of an Ornstein-Uhlenbeck process constrained to be positive. Therefore, we define $\tilde{k}_n$ as

$$\tilde{k}_n = \tilde{k}_{n-1} + \theta\,(\alpha - \tilde{k}_{n-1}) + \sigma\eta_n, \tag{5}$$

for suitable parameters (θ, α, σ), an initial condition $\tilde{k}_0$, and where $(\eta_n)_n$ is a sequence of i.i.d. standard normal random variables. Finally, for all $n \geq 1$, we set[6]

$$k_n = \max\{\tilde{k}_n, 0\}. \tag{6}$$

[5]We stress that the market maker also operates on a fast scale: he adjusts prices on the fast intra-day time scale and, eventually, makes the closing daily price available for time series analysis.

[6]The truncation to positive values is just one of the possible positive transformations of the Ornstein-Uhlenbeck process. See, for instance, Almgren (2012) for other specifications of positive transformations for mean-reverting stochastic processes representing liquidity on markets.

2.2 The artificial market at work: endogenous thresholds and intra-day returns

Each calendar date $n \in N$, the population I of agents trade on the market based on the liquidity prevailing that day (signalled by market depth k_n) and cost μ_n. In particular, at each intra-day time step $t \in \{1, \dots, \tau_n\}$, agents are asked to decide their trading strategy. Indeed, each agent $i \in I$ is either endowed with the asset or not. Suppose he does not own it, so that $\omega_i(t-1) = 0$ (a similar argument holds for the case of ownership). Therefore, he can either be a buyer or be idle. As described in Table 2, he reads the private signal ϵ_i and compares it with the computed threshold $T_i^b(t)$, where

$$T_i^b(t) = -\mathbb{E}^i[R_n(t)] + \mu_n = -k_n \left(\mathbb{E}^i \left[\frac{1}{I} \sum_{j=1}^{I} \omega_j(t) \right] - Q_n(t-1) \right) + \mu_n,$$

and $\mathbb{E}^i[\cdot]$ is the expectation with respect to the joint distribution of vector $\underline{\epsilon^{-i}} = (\epsilon_j)_{j \neq i}$. Note that this expectation is justified by the fact that ω_j depends on ϵ_j for all $j \neq i$ and that vector $\underline{\epsilon^{-i}}$ is not observed by the i-th agent. As said, the mechanism is strategic: each agent forms his expectation about all other agents' actions. Given payoffs as in Table 1, agents who perceive that they would be better off changing their strategy (from $\omega = 0$ to $\omega = 1$, or vice versa) will act accordingly. This will cause a change in $Q_n(t)$ and $R_n(t)$, so that other agents could also possibly decide to change their minds. Indeed, the procedure ends when all agents are happy with their actual strategy profile, meaning that the system has reached a Nash equilibrium $\underline{\omega}^*$. Each equilibrium vector $\underline{\omega}^*$ is in a one-to-one relationship with a vector of thresholds $\underline{T} = (T_i^b, T_i^s)_{i \in I}$.

In this respect, the vector of thresholds $\underline{T} = (T_i^b, T_i^s)_{i \in I}$ can be interpreted as a Nash equilibrium in an I-player game played at time step t. We now show the existence of such an equilibrium and a possible way to identify it explicitly as an iteration of a suitable *best-response map*. For reasons of clarity, as time step t is fixed, in the rest of this section we omit the time index. First of all, note that, for any buyer (and, similarly, for the seller),

$$\omega_i = \mathbb{I}_{\{\epsilon_i > T_i^b\}}.$$

The binary action ω_i translates into the real threshold $T_i^b \in \mathbb{R} \cup \{-\infty, +\infty\} := \bar{\mathbb{R}}$. The vector of actions is therefore $\underline{T} = (T_i^b, T_i^s)_{i \in I} \in \bar{\mathbb{R}}^{2N}$. Finally, call $\underline{T}^{-i,*}$, the $(I-1)$-dimensional vector formed by the optimal strategies excluding agent i. It can be shown that map $T_i(\underline{T}^{-i,*})$ is continuous (see Dai Pra et al. (2013) for more details). Being the best-response map $\underline{T} \mapsto \underline{T}^*$ continuous on the compact and convex domain $\bar{\mathbb{R}}^{2N}$, it admits at least one fixed point. Thus, we conclude that at least one Nash equilibrium exists. We have basically proved the following proposition.

Proposition 1. *For each date $n \in N$ and each intra-day time step $t \in \{1, \dots, \tau_n\}$, at least one Nash equilibrium exists, and it can be represented by threshold vectors*

$$(\underline{T}^b(t), \underline{T}^s(t)).$$

This finding resembles previous models in the field of agent-based models for trading. For example, in Ghoulmie et al. (2005), probabilistic time-varying thresholds are based on past returns. The main novelty of our methodology is that the thresholds we derive are implicitly determined at the equilibrium by assuming that agents forecast the level of returns prevailing on the market in the next time step and act strategically in a game-theoretical setting.

Having proved that the thresholds are well-defined, agents play according to the rules as in Table 2. We now briefly address the mechanism under which the Nash equilibrium can be numerically computed. Suppose we are at date n and at the beginning of the intra-day time step t. Each agent is endowed with common information about the market outlook. In particular, k_n, μ_n, and $Q_n(t-1)$ (the actual demand) are known to all agents. Moreover, each agent i knows his actual strategy $\omega_i(t-1)$ and receives a new private signal $\epsilon_i(t)$, which will play a fundamental role in a possible revision of his strategy. In order to compute the Nash equilibrium, we start with the actual vector of actions $\underline{\omega}(t-1)$. We denote it by $\underline{s}(0)$, this being the starting point of a new iteration procedure of the best-response map. Now, by sequentially playing according to a randomly selected order, each agent decides his *best action*. In doing so, he assumes that all other agents take their preferred action; this may imply a change in one entry of the state vector: $\underline{s}(0) \mapsto \underline{s}(1)$. The best-response iteration stops when no agent is willing to revise his own action, meaning that a Nash equilibrium $\underline{\omega}^*$ is determined. As seen, this equilibrium is characterized, equivalently, by a vector $(\underline{T}^b, \underline{T}^s)$ of real thresholds. Moreover, the related demand $Q_n(t) = \sum_i \omega_i^*/I$ is revealed. Returns $R_n(t)$ are realized on the intra-day market, and daily closing prices are computed. In the next section, we analyze the statistical properties of daily returns, prices, and market activity.[7]

3 Results

In this section, we present the results of the model. In particular, the next subsection outlines the procedure used to determine a set of parameters that give rise to several stylized facts, which are examined in Subsection 3.2. Finally, in the last part of the section, non-null transaction costs are introduced, and we describe their effects on returns and on market activity.

3.1 Parameters

Some parameters are related to the market, whereas others reflect traders' characteristics.

- Market characteristics. The only parameter related to the asset is the initial price p_0; without loss of generality, we set $p_0 = 1$. Concerning the time span,

[7]The Matlab code used to numerically simulate our returns, including the iteration procedure of the best response map, can be downloaded at `https://drive.google.com/open?id=1Gx-UpwslWZPpPvabaARFtj5YSWpw4weg`. The simulations produced by this code are exactly the ones that are statistically analyzed in Section 3.2.

we analyze time series formed by $N = 360$ calendar dates. As seen, the number of intra-day periods at dates $n \in N$ are i.i.d. random variables $(\tau_n)_{n\in N}$ with distribution $Pois(\bar{\tau})$. We fix $\bar{\tau} = 30$.[8]

- Traders characteristics. As discussed, at each trading date n, agents receive random signals $\epsilon_i(t)$, for $i = 1, \ldots, I$ and $t = 1, \ldots, \tau_n$. Signals are independent in time and across agents, and are distributed according to a centered logistic distribution with parameter $\beta = 10$. Concerning the number of traders, we assume that $I = 1000$ agents are active on the market.[9] Among them, a certain proportion $Q_n(0)$ is endowed with the asset. We assume that $Q_n(0)$ is random and, being a fraction of $I = 1000$, we model it as a Beta distribution $Beta(b_1, b_2)$, where $b_1 = b_2 = 4$ (truncating it at the second digit). The values of the parameters have been chosen to produce a bell-shaped symmetric distribution on $(0, 1)$, centered at 0.5.

The parameters related to the market depth (liquidity) signal, i.e. α, σ, θ, as defined in equation (5), are yet to be determined. They have been simply (but crudely) adjusted in order to generate some *stylized facts* of real financial markets. We use a quadratic fit measure to match some features of daily returns. In particular, we aimed to get the empirical average, standard deviation, and kurtosis of returns to be close to 0, 0.01, and 7, respectively.[10] In addition, we aspired to have approximately null first-lag autocorrelation of returns and first-lag autocorrelation of absolute returns close to 0.2. The reference values are reported for convenience in Table 3.

av_{ref}	sd_{ref}	ku_{ref}	$coret_{ref}$	$coabs_{ref}$
0	0.01	7	0	0.2

Table 3: Reference values for statistics used for the fit function.

To simplify the analysis and to reduce the non-trivial numerical complexity, we initially proceed by setting $\theta = 1$; in this way, we basically disregard the *memory* of the mean-reverting exogenous market depth process, and concentrate on the moments of the distribution of returns. Note that when $\theta = 1$, $(\tilde{k}_n)_n$ turns out to be a sequence of independent random variables distributed according to $\mathcal{N}(\alpha, \sigma)$. We

[8]The level of $\bar{\tau}$ is not crucial, provided that it is large enough to reach a stable value for intra-day returns. A value of $\bar{\tau}$ that is too low would not let the intra-day market stabilize on an *equilibrium* value for returns.

[9]The values of β and I are standard when studying the behavior of large populations using binary choice models. See, for instance, Phan et al. (2003) and Nadal et al. (2005), where the same value of β and a similar value for the number of agents is used.

[10]We used the `EuStockMarkets` dataset, which is available in R Core Team (2017) and contains the time series of the DAX, the SMI, the CAC and the FTSE from the beginning of 1991 to the end of 1998 (1860 daily observations for each index) to compute summary statistics of the returns. For instance, standard deviations and kurtosis are 0.010, 0.009, 0.011, 0.008 and 9.28, 8.74, 5.39, 5.64 for the four time series, respectively. Similar figures are reported as examples for stocks and indices in Campbell et al. (1997), even though there is obviously considerable variability among different assets or financial acitvities. In Pagan (1996), the estimate of autocorrelation at lag 1 of squared returns for US stocks is 0.189.

minimize the following reduced form fit function

$$F(\alpha, \sigma|\theta = 1) = (av - av_{ref})^2 + (sd - sd_{ref})^2 + (ku - ku_{ref})^2,$$

for a range of values of α and σ over a grid where $\alpha \in [0.02,\ 0.2]$ and $\sigma \in [0.02,\ 0.2]$. This analysis suggests that $\hat{\alpha} = 0.11$ and $\hat{\sigma} = 0.03$.[11]

Given the lack of memory in the process $\tilde{k}_n$, as expected, time series of returns do not show autocorrelation of absolute returns. We then keep $\hat{\alpha} = 0.11$ and $\hat{\sigma} = 0.03$, and look for the value of θ that approximately minimizes over θ the full fit function

$$F(\theta) = (av - av_{ref})^2 + (sd - sd_{ref})^2 + (ku - ku_{ref})^2 + (coret - coret_{ref})^2 + (coabs - coabs_{ref})^2, \tag{7}$$

where $\theta \in (0, 1]$. We find $\hat{\theta} = 0.4$. In Table 4, we collect all values of the parameters described above and implemented in the simulations.

α	σ	θ	μ	p_0	N	$\bar{\tau}$	β	I	b_1	b_2
0.11	0.03	0.4	0	1	360	30	10	1000	4	4

Table 4: Parameter values used in the simulations.

3.2 Returns, prices, and activity: analysis and discussion

In this section, we discuss the main statistical properties of the time series of daily returns generated by the model using the values of parameters as in Table 4. Figure 1

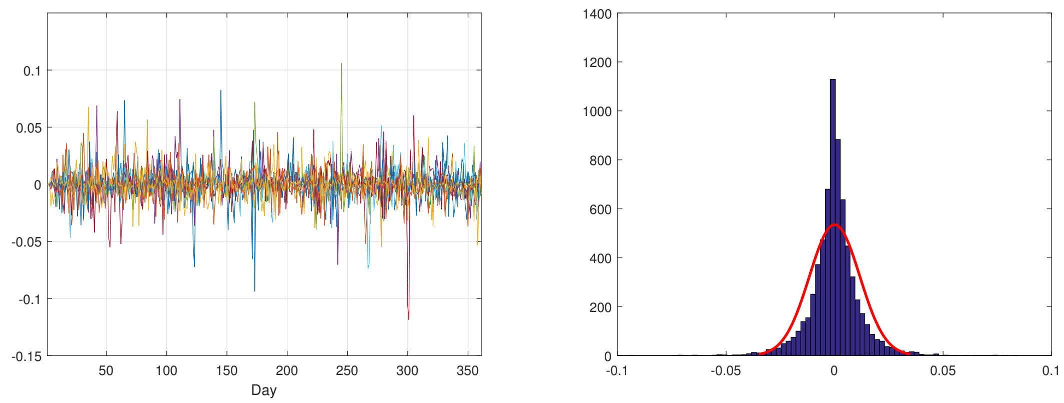

Figure 1: Time series (left) and histogram (right) of daily returns for the 19 simulations produced. Parameters are set as in Table 4.

[11] To run the numerical fit, we fix a seed for the generation of the random signals of the Ornstein-Uhlenbeck process. We then use a 20×20 grid of values for μ and σ. This calibration operation required about 30 hours of machine time on a Core i7-6700 processor. The same methodology was implemented in a second round for the calibration of θ, where we used a grid of ten values ranging from 0.1 to 1. Note that implementation of a Monte Carlo experiment with $M = 20$ trajectories would require approximately 20 days. A statistical robustness check of this result follows from the analysis of time series produced by the model relying on those values of the parameters.

shows how returns are usually close to zero, with some exceptions where pronounced and visible peaks are present on both the positive and negative side. In Table 5, we report the median values for the main statistics of the time series of returns.[12]

av	*sd*	*ku*	*coret*	*coabs*
−0.000007	1.01%	7.374	−0.063	0.241

Table 5: Summary statistics for time series of daily returns.

The average of the returns is basically null. The standard deviation, about 1%, is a reasonable value for many risky assets and traded financial activities. The distribution of the returns has substantial kurtosis, with the median exceeding 7. Moreover, some persistence of volatility is demonstrated by the 1-lag autocorrelation of absolute returns, 0.241, and by autocorrelograms reported in Figure 2. In this figure, we see an absence of autocorrelations of returns (left panel), but significant autocorrelations of absolute returns for a few lags (right panel).

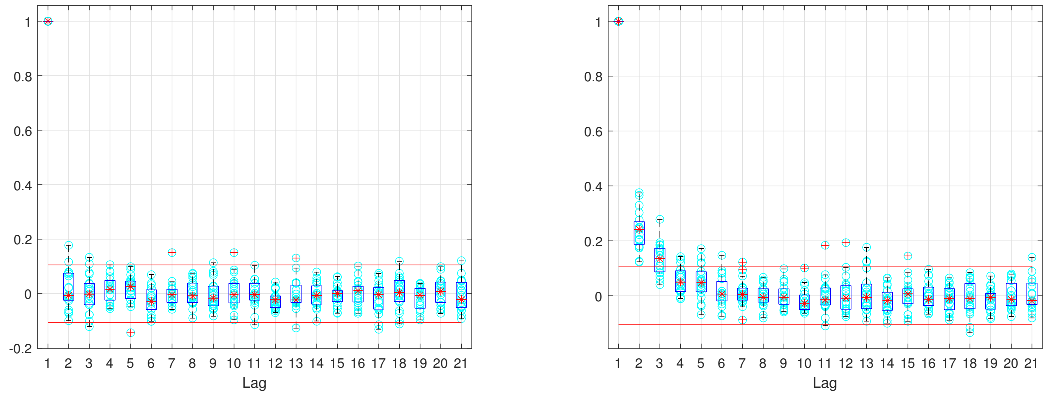

Figure 2: Autocorrelograms of daily returns (left panel) and absolute returns (right panel). As usual, the boxplots depict the median value, the inter-quartile range (the box); and outlying values marked with circles.

Simulated returns are fat-tailed and approximately power-law distributed, as seen in Figure 3, where we plot on log-log scales probability $P(r_n > R)$ in a representative case. We estimated power exponents for the 19 simulated series, see Gillespie (2015), and obtained results in the range $[2.50, 7.23]$. This is in very good agreement with Cont (2001), where it is claimed that the tail index is higher than 2 and less than 5 for many financial data (to be precise, 15 out of 19 simulated series have tail exponents between 2 and 5). Note that no attempt was made to obtain realistic tail exponents in the calibration of the parameters, and this unexpected outcome appears to validate the model's overall goodness of fit.[13]

It is also interesting to contrast the time series of daily returns and market depth k_n. We choose Simulation 9, which shows the highest abnormal return ($r_n \approx 10.6\%$

[12] Figures in Table 5 are computed as the median values extracted from the 19 simulations.

[13] We thank an anonymous referee for his or her useful remarks on tail exponents.

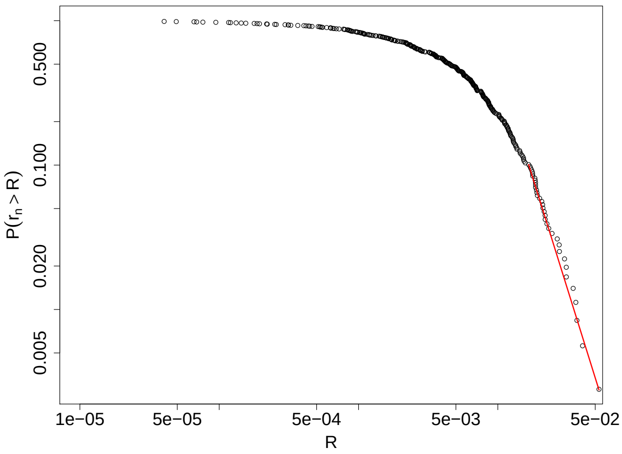

Figure 3: Log-log plot of $P(r_t > R)$ for Simulation 19. The fitted red line on the right tail has a -4.09 slope.

at calendar date $n = 244$). In this case, as is immediately apparent in Figure 4, we also see a peak in liquidity at the same date. We can interpret this fact as follows: agents on the market leverage on the liquidity shock to obtain abnormal returns on the fast-trading time scale. Note that the occasional presence of a large value of the market depth is not automatically transformed into excess returns: at date $n = 223$ of the same trajectory, we see a peak in k, which is not immediately realized on the market as an abnormal return. This suggests that our model, by mimicking a strategic mechanism, shows different possible reactions of traders to exogenous signals. This intuition is reinforced by statistically testing for the presence of linear correlation between returns and market depth signals. The hypothesis of null correlation cannot be rejected for all the simulations. Conversely, the hypothesis of null correlation between absolute returns and market depth must be rejected. In particular, considering the 19 simulations, we find a median correlation between absolute returns and market depth of 0.373 (maximum and minimum values are 0.440 and 0.293, respectively).

In Figure 5, we plot time series of daily prices. To improve legibility, we represent the ten odd trajectories generated by our simulations. On visual inspection, the time series resemble geometric Brownian motions with jumps. In particular, we recognize an upward jump at date $n = 244$ for Simulation 9 (the green line in the plot).

Another significant aggregate measure describing one important dimension of the market is activity. Using $A(t)$ to denote the proportion of agents trading in a certain intra-day trading period, it turns out that

$$A_n(t) = \frac{1}{I}\sum_{i=1}^{I}\omega_i(t)(1-\omega_i(t-1)) + \frac{1}{I}\sum_{i=1}^{I}\omega_i(t-1)(1-\omega_i(t)), \text{ for } t \in \{1,\ldots,\tau_n\}. \quad (8)$$

Note that $A_n(t)$ proxies trading volumes at time step t: in our stylized market, $A_n(t)$ coincides with volumes, since agents can only trade one share of the asset per period. In order to compare market activity with other figures, we compute a_n, an

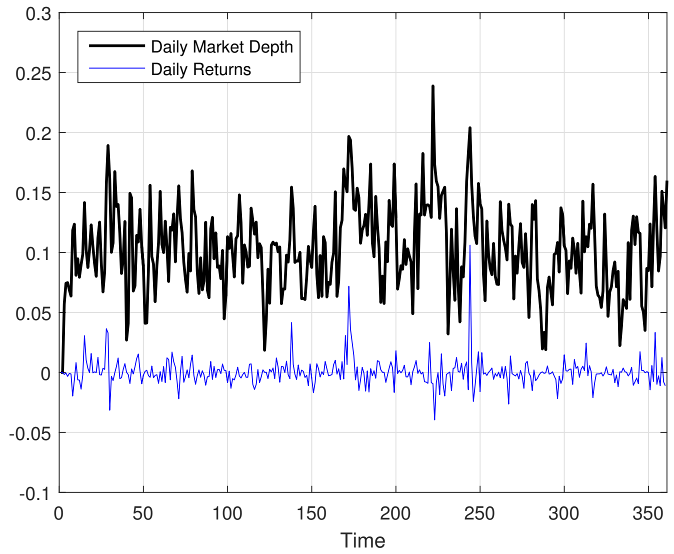

Figure 4: Comparison between daily returns (blue bottom line) and the exogenous signal of market depth (black upper line) for Simulation 9.

average daily activity, defined as

$$a_n = \frac{1}{\tau_n} \sum_{t=1}^{\tau_n} A_n(t).$$

In Figure 6, we plot a_n together with the daily returns r_n for Simulation 9. As expected, the daily returns have peaks corresponding to dates involving a high level of activity. The median value of daily activity across the 19 simulations is 0.5018, meaning that, without any taxation, we expect approximately half of the agents to trade one share of the asset at each intra-day period. We have also computed a proxy for the correlation between market activity and market volatility. The median value for the 19 simulations is 0.284.[14]

3.3 The impact of transaction costs on high-frequency trading

We have described a microfounded market for fast trading, where a large population of agents strategically interact to forecast returns and take advantage of high-frequency trends in prices. We have seen that our market mechanism is capable of reproducing the most significant stylized facts of real time series of returns and prices. It remains to discuss the role of transaction costs and whether it is possible to control high-frequency activities without affecting market performance.

To this end, we distinguish between two different approaches: a *passive* policy, where the market maker simply chooses a predetermined level $\mu > 0$ to reduce

[14] As a proxy for volatility of daily returns, we have used the standard deviation of intra-day returns.

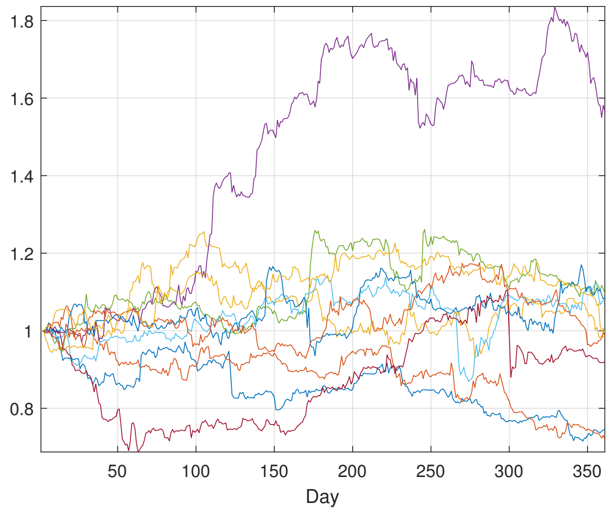

Figure 5: Time series of daily prices, for the simulations with odd numbers (from 1 to 19).

fluctuations of returns and excessive market activity. The advantage of this approach is firmly rooted in its simplicity. However, such a simplistic scheme can harm the market in periods of low activity. The second approach is a *proactive* policy, where the market maker reacts day by day to the signal about market liquidity and sets time-varying transaction costs. To keep the analysis simple, we analyze a policy that resembles the idea proposed by Spahn in one of his famous contributions, Spahn (1995). In particular, we propose a two-tier policy: the transaction cost $\mu_n > 0$ is *switched on* only at dates when market depth k_n exceeds a certain upper threshold k and, therefore, the risk of abnormal returns increases.

Approach 1: Constant transaction costs

As mentioned above, the first approach deals with a fixed non-zero transaction cost. To show the effects of μ on returns and activity, we consider two different scenarios. Scenario 1 with $\mu = 0.005$ and Scenario 2 with $\mu = 0.05$. The results are summarized, in terms of stylized facts, in Table 6 and in Figure 7.

Scenario	Parameters	*av*	*sd*	*ku*	*coret*	*coabs*	Activity
Scen. 1	$\mu = 0.005$	0.0000383	1.02%	7.962	0.011	0.221	0.5009
Scen. 2	$\mu = 0.05$	0.0000277	0.50%	8.581	−0.008	0.200	0.4691

Table 6: Summary statistics for time series of daily returns and activity with $\mu = 0.005$ and $\mu = 0.05$.

The introduction of a low-level cost basically has no significant consequences on the market. In contrast, a high level of transaction costs has a considerable effect on the variability of returns but, on the other hand, it harms activity: in this simulation, we see that the loss in activity is about 7% compared to the baseline scenario with

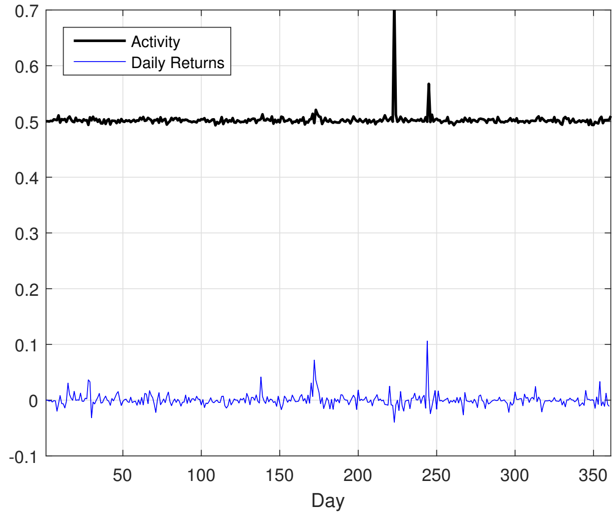

Figure 6: Time series of daily returns and market activity for Simulation 9.

$\mu = 0$, where the median value for activity a_n is 0.5018. Moreover, since returns are generally flat, few abnormal returns make the kurtosis higher compared to the $\mu = 0$ baseline case (see Table 5).

Approach 2: Time-varying transaction cost

The second approach, referred to previously as the proactive policy, is more sophisticated and considers the possibility of introducing a transaction cost only in case of a liquidity shock on the market (signalled by a high value of k_n). In the discussion of Figure 4, we point out that liquidity shocks may have a great impact on returns, inflating statistics such as standard deviations and kurtosis: the sudden reaction of fast traders to these shocks amplifies the effect of bad signals, affecting market outcomes. To control for this issue, we analyze a cost policy that resembles some ideas of the transaction tax proposed by Spahn. Indeed, we consider a two-tier policy such that $\mu_n > 0$, but only when k_n is large enough. We have, for all $n \in N$,

$$\mu_n = \mu \cdot \mathbb{I}_{\{k_n \geq k\}}, \tag{9}$$

where $\mu > 0$ is a predetermined cost and k is a suitable threshold level for the market depth parameter k_n. We must then determine μ and k. On the one hand, sustained activity is needed to keep the market vital; on the other hand, overly high levels of liveliness make the returns too volatile. The market maker is likely to strike a balance between these two conflicting goals. Again, we consider two different scenarios: Scenario 3, where $k = 0.1$ and $\mu = 0.01$; and Scenario 4, where $k = 0.1$ and $\mu = 0.05$. Statistics are reported in Table 7 and results displayed in Figure 8. Scenario 3, characterized by $\mu = 0.01$, shows no significant difference compared to the baseline case with $\mu = 0$. Scenario 4, where a rather significant cost ($\mu = 0.05$) is imposed only if market depth is above threshold $k = 0.1$, turns out to be more interesting. In this case, variability (standard deviation) is much lower than in the baseline scenario of frictionless markets, and is in line with the case of a flat cost $\mu = 0.05$ (see Scenario 2 in Table 6). Interestingly, activity is not harmed as much as in the case of Scenario 2: activity is now reduced only by 3.9% compared to the baseline case without frictions.

Scenario	Parameters	*av*	*sd*	*ku*	*coret*	*coabs*	Activity
Scen. 3	$k = 0.1$ & $\mu = 0.01$	0.0000459	0.97%	7.945	0.019	0.201	0.5003
Scen. 4	$k = 0.1$ & $\mu = 0.05$	−0.000494	0.55%	7.306	−0.010	0.119	0.4822

Table 7: Summary statistics for time series of daily returns and activity under a Spahn-style policy with $k = 0.1$ and under two scenarios for the level of cost $\mu = 0.01$ and $\mu = 0.05$, respectively.

4 Conclusions

In this paper, we discuss an artificial market where actions unfold according to two time scales: in the model, many high-frequency traders interact using simple threshold rules, strategically attempting to maximize their short-term returns and taking into account the whole set of orders produced by other agents to reach a local Nash equilibrium; on a much slower scale, (inverse) liquidity in the market evolves following an autocorrelated exogenous process, and a market maker sets the closing daily price and can adjust transaction costs to reduce, say, outlying returns and curb excess kurtosis, while maintaining sustained activity in the marketplace.

The baseline model with no transaction costs, after some calibration aimed to attain a host of relevant stylized facts, is indeed able to reproduce several significant statistical regularities of realistic financial time series, including fat-tailed returns, sizeable excess kurtosis, no linear predictability of returns, and some volatility clustering. The model thus suggests that fast strategic interaction at the intra-day level can amplify liquidity-related shocks and trigger streams of orders leading to outlying returns, even in the presence of a slow and conservative market mechanism.

Quite interestingly, a liquidity stochastic process with memory turns out to be filtered by the artificial market and is morphed in a white sequence of uncorrelated returns, which are no longer normally distributed and exhibit volatility bursts. While we believe that the results can chiefly be ascribed to the effects of the fierce and fast strategic interaction occurring during a trading session, additional research would be needed to endogenize the prevailing liquidity and clarify the role of some key parameters: θ, for example, appears to play an important role in obtaining some stylized facts that are generated only if $\theta < 1$, thus inducing some (small) predictability in the sequence of daily market depth k_n.

Finally, we discuss the impact of a transaction cost on daily returns, prices, and market activity. In particular, we propose two alternative mechanisms: (i) a fixed cost to be paid at each transaction on the realized returns, irrespective of the exogenous signal of market depth; (ii) a variable cost to be paid only when the market depth exceeds a predetermined threshold. The model suggests that the latter mechanism, referred to as the *proactive* policy, may help in preventing the amplifying effects of fast trading over the liquidity-related signal. On the other hand, at dates when the signal is ordinary (i.e. below the activation threshold), the transaction cost is set to zero so that activity is not detrimentally harmed. In this perspective, our proactive policy resembles the two-tier taxation proposed in Spahn (1995).

Our results further reinforce a pattern that has been gaining momentum in recent studies, and stress the still-overlooked fact that substantial agent heterogeneity is needed if credible returns are to be generated (in our model, orders are induced by heterogenous thresholds and signals that are independently and repeatedly drawn while trading). At the same time, the microstructure of the model is quite simplistic; more detailed exchange structures than a simple market maker may improve the quality and intensity of the stylized facts detected in the simulated data.

5 Acknowledgments

We are indebted to two anonymous referees for their helpful and constructive comments. The authors acknowledge the financial support of Ca' Foscari University of Venice under the grant *"Interactions in complex economic systems: innovation, contagion and crises"*. We are also grateful to Antonella Basso for her support in sharing computer multi-core resources that proved to be invaluable for running our simulations. We received useful remarks from participants in the WEHIA 2017 conference held at Università Cattolica di Milano and seminars held at the University of Technology Sydney, Ca' Foscari University of Venice and at the Centre d'Economie de la Sorbonne (CES), but are entirely responsible for all remaining errors.

References

Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics*, 3(1):163–181.

Brock, W. A. and Durlauf, S. N. (2001). Discrete choice with social interactions. *The Review of Economic Studies*, 68(2):235–260.

Budish, E., Cramton, P., and Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621.

Campbell, J. Y., Lo, A. W.-C., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.

Chiarella, C., He, X.-Z., Shi, L., and Wei, L. (2017). A behavioural model of investor sentiment in limit order markets. *Quantitative Finance*, 17(1):71–86.

Chiarella, C., Iori, G., and Perelló, J. (2009). The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Economic Dynamics and Control*, 33(3):525–537.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236.

Dai Pra, P., Sartori, E., and Tolotti, M. (2013). Strategic interaction in trend-driven dynamics. *Journal of Statistical Physics*, 152(4):724–741.

Fontini, F., Sartori, E., and Tolotti, M. (2016). Are transaction taxes a cause of financial instability? *Physica A: Statistical Mechanics and its Applications*, 450:57–70.

Ghoulmie, F., Cont, R., and Nadal, J. (2005). Heterogeneity and feedback in an agent-based market model. *Journal of Physics: condensed matter*, 17:1259–1268.

Gillespie, C. S. (2015). Fitting heavy tailed distributions: The poweRlaw package. *Journal of Statistical Software*, 64(2):1–16.

Granovetter, M. (1978). Threshold models of collective behavior. *The American Journal of Sociology*, 83(6):1420–1443.

Hasbrouck, J. and Saar, G. (2009). Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12(2):143–172.

Hasbrouck, J. and Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*, 16(4):646–679.

Kaldor, N. (1961). Capital accumulation and economic growth. In Lutz, F. A. and Hague, D. C., editors, *The Theory of Capital*. St. Martin's Press.

Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136.

LeBaron, B. and Yamamoto, R. (2008). The impact of imitation on long memory in an order-driven market. *Eastern Economic Journal*, 34(4):504–517.

Lux, T. (2009). Stochastic behavioral asset-pricing models and the stylized facts. In Hens, T. and Schenk-Hoppe, K. R., editors, *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance, pages 161–215. North-Holland, San Diego.

Lux, T. (2016). Applications of statistical physics methods in economics: Current state and perspectives. *The European Physical Journal Special Topics*, 225(17):3255–3259.

Maslov, S. (2000). Simple model of a limit order-driven market. *Physica A: Statistical Mechanics and its Applications*, 278(3-4):571–578.

Nadal, J.-P., Phan, D., Gordon, M. B., and Vannimenus, J. (2005). Multiple equilibria in a monopoly market with heterogeneous agents and externalities. *Quantitative Finance*, 5(6):557–568.

Pagan, A. (1996). The econometrics of financial markets. *Journal of Empirical Finance*, 3(1):15–102.

Phan, D., Pajot, S., and Nadal, J.-P. (2003). The monopolist's market with discrete choices and network externality revisited: Small-worlds, phase transition and avalanches in an ace framework. In *Ninth Annual Meeting of the Society of Computational Economics, University of Washington, Seattle, USA*.

R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Spahn, P. B. (1995). International financial flows and transactions taxes: survey and options.

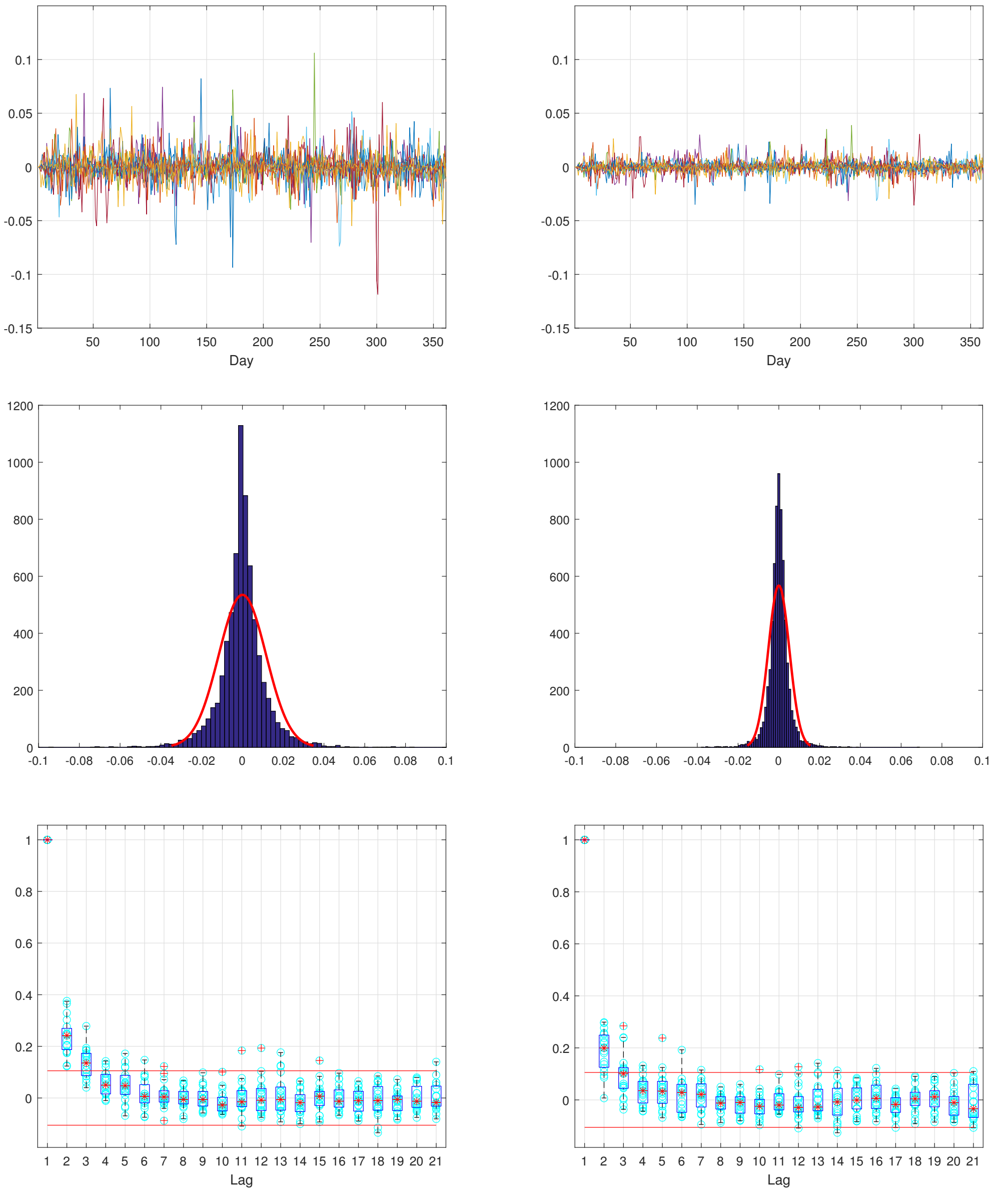

Figure 7: Time series of daily returns (top), histogram of returns (center), correlograms of absolute returns (bottom) under Scenario 1 (left panels) and Scenario 2 (right panels).

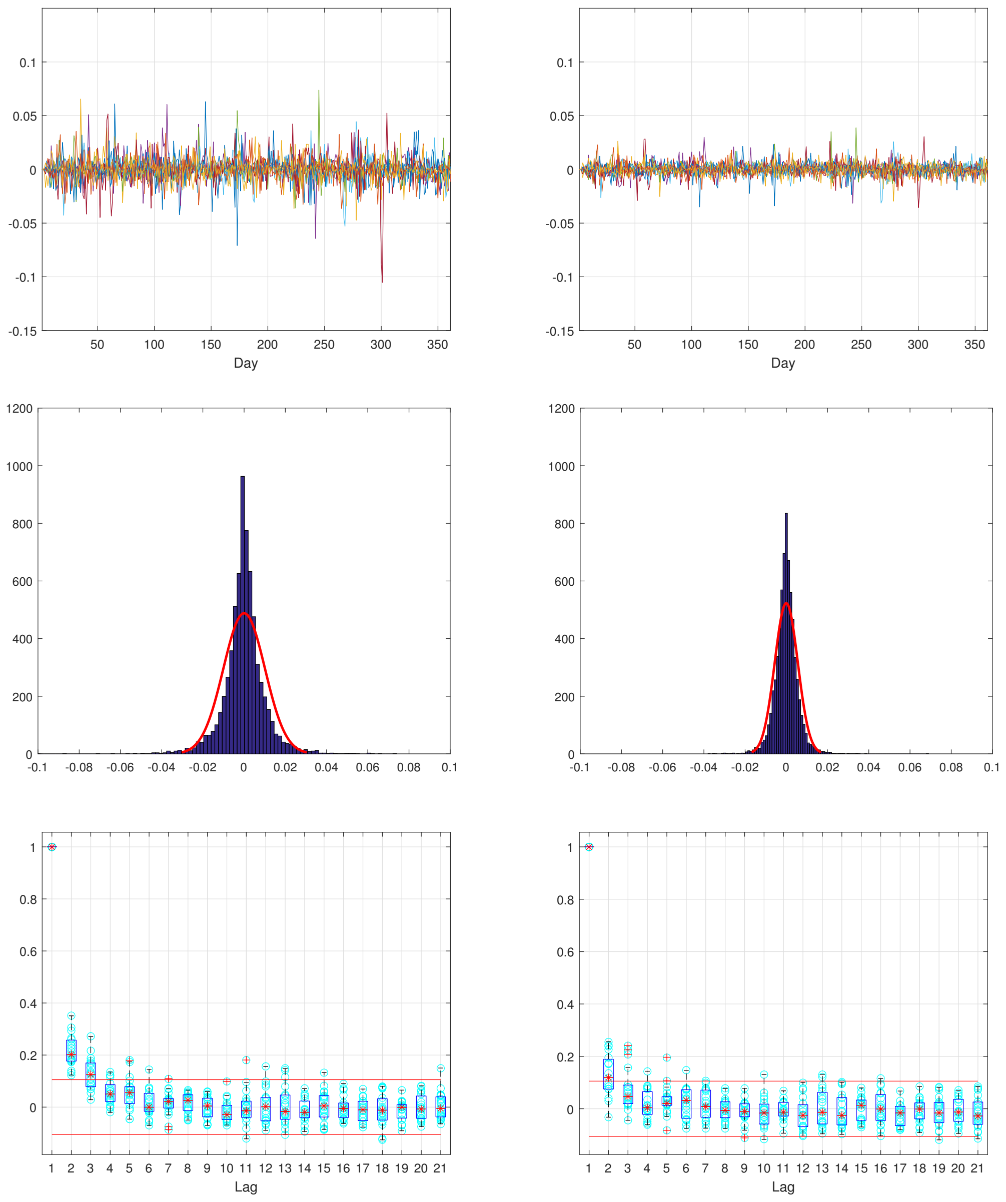

Figure 8: Time series of daily returns (top), histogram of returns (center) and correlograms (bottom) under Scenario 3 (left panels) and Scenario 4 (right panels).

Click here to download Data Deposition Information INCLUDING WEBLINK https://drive.google.com/open?id=1Gx-UpwslWZPpPvabaARFtj5YSWpw4weg