

Fundamentalists clashing over the book: a study of order-driven stock markets

Marco LiCalzi¹ and Paolo Pellizzari

Department of Applied Mathematics, University of Venice,
Dorsoduro 3825/e, 30123 Venezia, Italy

E-mail: licalzi@unive.it and paolop@unive.it

Received 20 July 2002, in final form 30 May 2003

Published

Online at stacks.iop.org/Quant/3

Abstract

Agent-based models of market dynamics must strike a compromise between the structural assumptions that represent the trading mechanism and the behavioural assumptions that describe the rules by which traders make their decisions. We present a structurally detailed model of an order-driven stock market and show that a minimal set of behavioural assumptions suffices to generate a leptokurtic distribution of short-term log-returns. This result supports the conjecture that the emergence of some statistical properties of financial time series is due to the microstructure of stock markets.

1. Introduction

The dynamics of a stock market depends on the interaction between the trading mechanism and the behaviour of the participants. The trading mechanism defines the rules of the market, which specify how orders are placed and how the price changes. The behaviour of the participants is the outcome of their trading strategies, which include how they form expectations or interpret signals.

An agent-based model of the stock market makes *structural* assumptions that represent the trading mechanism and *behavioural* assumptions that describe the rules by which traders take their decisions. The models in the literature strike different compromises between the accuracy of their structural and behavioural assumptions. We believe that the first generation of agent-based simulations of stock markets has explored a very rich set of behavioural assumptions, but has paid comparatively little attention to structural assumptions.

This neglect of structural assumptions in agent-based modelling has a few notable exceptions. Stigler (1964, p 125), for instance, sharply criticized a report prepared for the Securities Exchange Commission by pointing out that it was based on the structural assumption that ‘smoothness of price

movements is the sign of an efficient market’. By simulating a random stream of tenders, he showed that this assumption was unwarranted. Gode and Sunder (1993) established the allocative efficiency of a double auction regardless of any specific behavioural assumptions by comparing the experimental performance of human traders with the simulated choices of zero-intelligence traders.

The purpose of this paper is to present a structurally detailed agent-based model of an order-driven stock market and show that a minimal set of behavioural assumptions suffices to generate a leptokurtic distribution of short-term log-returns. While leptokurtosis is only one of several well known ‘stylized facts’ about financial time series, this finding backs up the conjecture that the emergence of some of the statistical properties of order-driven markets is mostly due to their microstructure.

With respect to one important point, this viewpoint is not new. It has been advocated by several recent studies in the physics literature about the statistical properties of order-driven markets; see among others Bouchaud *et al* (2002) and Daniels *et al* (2002). Their approach focuses on modelling the buying and selling order flows affecting the book as random processes and studies the emerging statistical properties associated with order-driven markets. While being very parsimonious, models

¹ Author to whom any correspondence should be addressed.

in this family exhibit impressive descriptive power and are able to generate falsifiable predictions. See in particular Smith *et al* (2002), which also provides methodological foundations and a short but enlightening discussion of prior related work.

On the other hand, by electing order flows as its starting point, this approach avoids the explicit modelling of agents' behaviour. It is plausible to expect that most combinations of reasonable behaviours should lead to sufficiently similar order flows, but giving up agent-based modelling makes it impossible to address this issue. Chiarella and Iori (2002) makes an important step towards blending a statistical approach with agent-based modelling. However, their behavioural assumptions are so rich that the impact of the microstructure remains hard to assess. Our paper aims to provide a stricter agent-based test of the relevance of structural assumptions.

This test is based on a model of an order-driven stock market that pays due attention to its structural features. Based on this model, we show that a minimal set of behavioural assumptions suffice to generate a bell-shaped but remarkably leptokurtic distribution of short-term log-returns. This leptokurtosis combines deviations from normality both around and far away from the mode of a bell-shaped density, which in the literature are conventionally known as 'narrow peak' and 'fat tails'. The minimality of the behavioural assumptions supports the conjecture that these two common statistical properties of financial time series stem from the market microstructure rather than from sophisticated behavioural assumptions.

1.1. Background

Since Beja and Goldman (1980), the dynamic behaviour of stock prices has often been explained as the outcome of the interaction of a few archetypal trading strategies. A common approach is to simulate a one-asset economy with three types of agents: fundamentalists, trend chasers and a market maker; see Day and Huang (1990). The fundamentalists agree on a long-term value v of the asset; they buy shares when the price p of the stock is below v and sell them otherwise. Trend chasers accumulate shares when p moves up and unload them when it goes down. The market maker is a specialist who matches excess demand out of his inventory and equilibrates the market.

The intuition supported in this approach is well known. Different families of trading strategies endow the dynamics of price formation with negative and positive feedbacks, whose interplay generates time series that are consistent with many qualitative and statistical properties of the price trajectories observed in real markets, such as sudden and unpredictable switchings between bullish and bearish regimes, short-term correlations, fat tails or clustered volatility; see for instance Lux (1998).

This (admittedly stylized) approach has at least two well known limitations, which have sparked a new wave of research about order-driven markets with finite resources. First, there is no explicit trading mechanism. The market maker is risk neutral and endowed with unbounded liquidity, which is used to absorb excess demand and make trading always viable, regardless of its size. In each period, the market maker

adjusts the price to reduce the excess demand. Inspired by the metaphor of the Walrasian auctioneer, this price-adjusting rule fails to recognize adequately that in a real market trade occurs whenever two agents can match their requests at a given price. This makes it impossible to test the following plausible explanation for the narrow peak of returns: the trading mechanism makes prices 'stickier' by increasing the probability that orders cannot find a match; hence, returns cluster around the mean.

Second, there is rarely mention of the liquidity constraints. Trend chasers or fundamentalists can keep loading their portfolio regardless of the amount of money they have already spent or of the short selling they have incurred. While it is possible to argue that most fundamentalists are large financial institutions with copious budgets, this is less credible for trend chasers who can end up being 'sheared' by the market forces. Moreover, the absence of liquidity constraints for the market maker implies that rationing never occurs and thus abnormally large orders must cause significant swings in price, which might be responsible for the fat tails. However, since rationing does occur in real markets, we do not know how much of the fat tails generated by this mechanism would survive in a more realistic model.

The (so far scarce) literature reporting agent-based simulations in order-driven markets initiated with Cohen *et al* (1978) and Mendelson (1982). They do not provide a reliable test for our conjecture, but offer some support for its plausibility. Bak *et al* (1997, p 439) simulates markets where agents place buy and sell orders; for the simple but unrealistic case of pure noise trading, it analytically proves that the 'anomalous scaling behaviour (observed in real markets) can arise as a consequence of the interactions between many agents'. This conclusion is backed up by the simulations of order-driven markets offered in Steiglitz and Shapiro (1998) and Chiarella and Iori (2002).

Similarly to us, Maslov (2000) claims that some price patterns may be a consequence of the trading mechanism and shows that the interaction of pure noise traders in order-driven markets leads to fat tails in the short-term log-returns density. This finding is confirmed in the Genoa market described in Raberto *et al* (2001); its abstract mentions the appearance of leptokurtosis, but narrow peaks are not explicitly discussed in the paper. Matassini and Franci (2001) obtain fat tails in order-driven markets where traders' strategies are based on a convex combination of value trading, herding and chartism.

For the purpose of testing our conjecture, the main limitation of these papers is that they assume either pure noise trading or the coexistence of a variety of trading strategies. Pure noise trading is not realistic, while the simultaneous existence of sharply defined classes of trading strategies is a strong behavioural assumption. For instance, consider how fundamentalists and trend chasers tend to stay on opposite sides of the market. Since only one side can gain from a transaction, it remains unclear why agents do not switch strategies; see Lux and Marchesi (2000).

One justification for clumping all the fundamentalists on the same side of the market is that, if they agree on the exact value of v , they cannot profitably trade among themselves;

see Milgrom and Stokey (1982). Hence, the presence of alternative strategies such as trend chasing or noise trading is necessary to make the market viable. However, perfect agreement about v requires extreme assumptions such as common knowledge of the informational partitions. Even financial analysts, who are supposed to look over fundamentals and are prone to herding effects that tend to cluster their predictions, rarely reach unanimous opinions; see Welch (2000).

It seems more realistic to have enough heterogeneity to make trade also possible within the class of fundamentalist trading strategies. This can be easily achieved by assuming that fundamentalists base their value trading on heterogeneous estimates of v . We interpret this heterogeneity as a clash of opinions between equally smart fundamentalists who use different estimation procedures or have access to different information, or between better and worse forecasters; see Arthur *et al* (1997).

A small degree of heterogeneity among fundamentalists suffices to achieve market viability even in the absence of other trading strategies. This provides a strong test for our conjecture. If an order-based trading mechanism produces fat tails and a narrow peak in a market where there are only fundamentalists, the emergence of these two effects must be due to the structural assumptions. A second and complementary test needs to show that the same effects arise in a market with zero-intelligence agents, whose trading choices satisfies minimal behavioural constraints that set them apart from pure noise traders or other irrational agents. We perform these two tests in a variety of environments to check for their robustness. The results strongly support the conjecture.

2. The model

2.1. Structural assumptions

We consider an economy with two assets: one bond and one stock. The bond pays a riskless yearly interest rate r . The stock has a risky price p , which depends on demand and supply, and pays no dividends. We assume that the interest earned on the bond is spent elsewhere, so that the economy is closed: the total amount of cash C and stock S in the economy is constant. However, since the number of active traders may fluctuate over time, the quantities of cash and stock available on the market are not constant. We assume that no new information is ever released and thus, except for the number of active traders, the fundamentals of the economy are essentially unchanged over time.

There are n (potentially active) traders, who enter or exit the market independently from each other. Upon entering the market, a trader i is endowed with a quantity c_i of cash and a quantity s_i of stock. An agent can use the money in his possession to buy stock, or he can sell the stock he owns in exchange for cash; but he is not allowed to short sell or to borrow money. Thus, a trader can place a buy order only if he has enough cash and a sell order only if he has enough stock; in short, he is budget constrained both on the buy and the sell side. Since the economy is closed, the gains and losses of active traders must offset each other.

The market is order driven. There is a book of orders that each trader can check at any time. Within his budget constraints, a trader can place market orders or limit orders for arbitrary quantities. A market order is filled completely if it finds enough capacity on the book, or partially otherwise. A limit order is stored in the book and executed (partially or completely) when it finds a match during the rest of the trading session. If a market order or a limit order are not filled completely, the agent is rationed. For realism, prices on the book must be quoted in ticks. The minimum tick allowed on the book is of the order of $1/1000$ of the stock price.

2.2. Behavioural assumptions

We assume that all agents are fundamentalists. This assumption combines two hypotheses: the first one concerns the trading strategy used (which is value based: buy low, sell high), the second one concerns the estimate of the fundamental value v . In a world where value-based traders are allowed to disagree, each trader i may have reasons to suspect that his estimate v_i is off the mark. For instance, if the price drifts away from v_i , trader i may be induced to revise v_i . Thus, even if the underlying strategy is value based, the estimate of v_i may be tinged with some form of chartism, opening up an interesting avenue of research. True to their name, however, our (die-hard) fundamentalists never revise their own estimates during their activity span. This strong assumption provides a more stringent test for our conjecture, because it rules out any indirect effects of price dynamics on the trading orders.

The value-based trading strategy used by our fundamentalists is the following. When an agent enters the market at time t , he has an investment horizon h_i and he wishes to maximize his gains over the time span $h_i - t$, which is his *activity period*. He formulates an estimate of the fundamental value v_i that he expects the stock will reach by time h_i . Since the bond has a riskless yearly rate of return r and investment in the stock is risky, trader i requires a yearly risk premium $\pi_i > 0$ to invest in the stock. One might include transaction costs in π_i or make the premium time dependent, but for simplicity we assume that it is constant per unit of time.

Let p the price of the stock at time t . Trader i is willing to buy stock at time t if he expects a return sufficiently higher than the return on the riskless bond; that is, if

$$\frac{v_i}{p} \geq 1 + (r + \pi_i)(h_i - t).$$

This implies that the highest bid price he is willing to offer at time t is

$$\beta_i(t) = \frac{v_i}{1 + (r + \pi_i)(h_i - t)}. \quad (1)$$

Similarly, trader i prefers to invest in the bond when the expected return from holding stock is lower than the riskless return achievable from investing in the bond. Hence, he is willing to sell stock if

$$\frac{v_i}{p} \leq 1 + r(h_i - t),$$

which gives

$$\alpha_i(t) = \frac{v_i}{1 + r(h_i - t)} \quad (2)$$

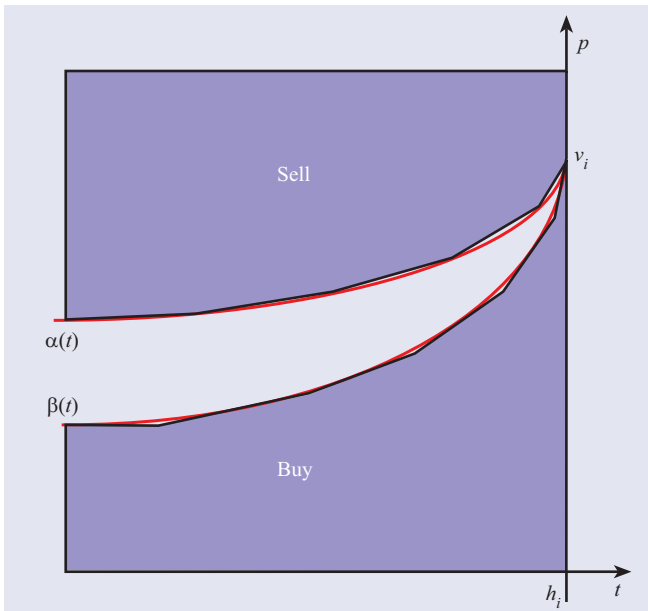


Figure 1. The bid–ask spread.

as the lowest ask price that he is willing to offer at time t . Except for its explicit time dependence, this behavioural rule is similar to the one suggested in Bak *et al* (1997) for the fundamentalists' trading decisions.

If the agent enters the market at time t_i , the two increasing functions $\beta_i(t)$ and $\alpha_i(t)$ are defined for t in $[t, h_i]$. During this interval of activity, and compatibly with his budget constraints, the agent respectively buys or sells the stock if the current price is below $\beta_i(t)$ or above $\alpha_i(t)$, and hold his position in between; see figure 1. This behavioural rule deserve two comments.

First, as $t \rightarrow h_i$, the spread between $\alpha_i(t)$ and $\beta_i(t)$ narrows down and the agent becomes more prone to 'taking positions'. This provides liquidity to the market. In particular, when $t = h_i$, the agent has reached his investment horizon and exits the market while keeping his endowment. The endowment of an old trader who has quit the market is transferred to a new trader who enters the market some time after, with a new investment horizon and a new estimate for the fundamental value of the stock. Since investment horizons are randomly distributed across traders, this implies a staggered entry of new traders which further increases the liquidity of the market. However, the primary liquidity driver is the sensitivity of agents' behavioural rules to time.

Second, it is plain that the existence of the bond plays no other role than providing the riskless rate of interest r used in the definition of $\alpha_i(t)$ and $\beta_i(t)$. Therefore, it could be effaced at no cost. However, it provides a useful metaphor to discuss changes in monetary policy, as we do below.

Let b and a be respectively the best bid price and the best ask price available when trader i checks the book. If $\beta_i(t) \geq a$, then he places a market order for purchasing stock at a price of a . If the supply available on the book at a price of a is sufficiently large, he only buys the stock he can afford; otherwise, he buys all the quantity available and then moves on to check the second best ask price, iterating the process until

the agent has no more cash or cannot find a n th level ask price on the book lower than $\beta_i(t)$. If this second condition occurs, the agent places a limit order at a price equal to the first tick below $\beta_i(t)$ for the maximum quantity of stock he can afford to pay for.

Symmetrically, if $\alpha_i(t) \leq b$ then he places a market order for selling stock at a price of b . If the demand available on the book at a price of b is sufficiently large, he sells only the stock he owns; otherwise, he fills the available demand and then moves on to check the second best bid price, iterating the process until the agent has no more stock or cannot find a n th level bid price on the book greater than $\alpha_i(t)$. If this second condition occurs, the agent places a limit order at a price equal to the first tick above $\alpha_i(t)$ for the quantity of stock he still owns. At the end of the process a trader might end up with two limit orders placed on different sides of the book, as he stands ready to profit from either an increase or a decrease in the price of the stock. Also, note that a trader finds his market orders more profitable than the corresponding limit orders.

We assume that trading takes place in sessions and, for convenience, we think of a trading session as 1 day. Since there are about 250 trading days in a year, we clock time in increments of $1/250$. Each of these increments corresponds to 1 day and one trading session. Within 1 day, all traders check the book and place orders asynchronously. The book matches order-satisfying price priority and, in case of equal prices, temporal priority.

Temporal priority implies a standard trade-off between efficacy and immediacy: a trader has both reasons to postpone placing his limit orders (in the hope of passing it as a market order) and to anticipate them (to limit the risk of being rationed). We solve this issue by randomizing with uniform probability the order in which traders check the book within 1 day and by clearing the book at the end of a trading session. The day after, we draw a new permutation over the set of active traders and let them all check the book and place again their orders. This implies that limit orders unexecuted by the end of a trading session are temporarily withdrawn and is consistent with a common practice on real stock markets such as NYSE, where most limit orders are 'for today only' rather than 'good until cancelled'.

Removing a limit order from the book between trading sessions protects a trader from the risk that the arrival of new information could force him to do a trade that is no longer convenient. The arrival of new information is ruled out in our model, where it is the time dependency of $\beta_i(t)$ and $\alpha_i(t)$ that requires a trader to update his limit prices. This effect is implicitly suggested in Bak *et al* (1997). In contrast, in the literature most agent-based simulations of an order book assume that unexecuted limit orders have longer lifetimes. This is irrelevant when both new information and time effects are irrelevant, as in the pure noise trading model of Maslov (2000). However, in models with mixed motives like Chiarella and Iori (2002) or Matassini and Franci (2001), longer expiration dates imply that traders with standing limit orders 'freeze' their motives for quite a while.

Table 1. The benchmark case.

	Parameters	Initialization
Global	n	= 3410
	r	= 0.02
	p	= 1000
Trader	c_i	= 2000 (first activation)
	s_i	= 1 (first activation)
	v_i	$\sim p[1 + U(-0.1, 0.1)]$
	π_i	$\sim U(0, 0.03)$
	h_i	$\sim t + \lceil \exp(1/250) \rceil$ days
	τ_i	$\sim h_i + \lceil \exp(1/250) \rceil$ days

3. Simulations with fundamentalists

3.1. Identification

A simulation run for our model requires the specification of three global parameters and a list of individual variables for each trader whenever he is activated. The global parameters are the number of potentially active traders n , the riskless rate of interest r and the initial stock price p . Upon activation, a trader i is characterized by his endowment of cash c_i and stock s_i and by a four-tuple $(v_i, \pi, h_i; \tau_i)$ of parameters, where v_i is his estimate of the fundamental value, π_i his risk premium and h_i is his activity period. When the active trader exits the market, he is replaced by a new trader after a period of time τ_i .

We assume that the initial endowment of cash and stock is identical across traders; successively, to keep constant the amount of stock and cash in the economy, a new trader inherits the same endowment of the trader that he replaces. For each trader which becomes active, the four-tuple $(v_i, \pi_i, h_i; \tau_i)$ is the outcome of a random realization. More precisely, for each of these four parameters we choose a probability distribution and set its value equal to an independent draw from that distribution. Hence, the set of agents' variables is described by a combination of endowments which depends on the previous history of the market and by a random realization of the parameters.

3.2. The benchmark case

We have performed very different simulations² for our model. However, we define for concreteness the benchmark case summarized in table 1, where U and \exp denote respectively the uniform and the exponential distributions. The symbol ' \sim ' denotes an independent draw from a probability distribution. The global parameters are set once and for all at time $t = 0$; the individual endowments c_i and s_i are set once upon activation and then inherited upon successive activations; the four individual parameters are reset each time a trader is activated.

We stress that this choice of parameters is loosely inspired by the empirical evidence, but it is not an attempt to calibrate the model to a specific set of real data. We neither believe nor claim that our model has reached the level of descriptive power that would warrant a serious calibration exercise. Our

simulations are not meant to replicate the price dynamics observed in real markets, but to show that a simple order-driven market can generate both fat tails and a narrow peak in the distribution of daily log-returns even if all the traders in the market are fundamentalists.

We comment briefly on the parameters. The number of potentially active traders is $n = 3410$. At the start of the simulation, exactly half of them are declared active. Successively, as old traders exit the market and new traders enter it, the number of active traders fluctuates around an average value of about 50% of the potentially active traders. The fluctuation spans a band thinner than $\pm 3.5\%$. As we explain below, $n = 3410$ is a convenient choice that allows a direct comparison in one of our subsequent tests of robustness. The yearly rate of return $r = 0.02$ is close to the current rate of inflation in the Euro zone, which is a reasonable lower bound for the nominal riskless rate of return.

At time $t = 0$, all traders who become active enter the market with an equal endowment of cash and stock. Later, as an agent inherits the cash and the stock left by an exited trader he is replacing, traders may have different endowments. The amount of initial endowments (combined with the initial distribution of v) ensures that the level of stock price is not constrained by the overall liquidity in the system: there is enough cash to support at least a stock price of 2000, while the actual prices are far below this level.

At time $t = 0$, the estimates of the fundamental value are uniformly distributed around 1000, with an offset of at most $\pm 4\%$. This choice induces the initial level of the stock price to hover around 1000, which improves the readability of the time series and allows us to set the size of a tick equal to 1. The initial offset is smaller in size than the 10% given in table 1 to cap the price oscillations occurring when the market is jumpstarted. Traders who become active later and enter the market at some time $t > 0$ are given an estimate v_i uniformly distributed around the last official³ price p_{t-1} with an offset of $\pm 10\%$. The risk premium π can range from 0 to 1.5 the rate of interest, in order to span a large set of possible attitudes to risk.

Consistent with the standard assumption that entries in and exits from the market follow (independent) Poisson processes, both activity and inactivity periods are exponentially distributed with an expected length of 250 trading days, equivalent to 1 year. In particular, at time $t = 0$, this assumption is used to initialize both the investment horizons of the traders declared active and the inactivity periods of the other traders. One might assume different expected lengths for the two periods or even make these a function of other parameters of the model. For simplicity, we assume that they are constant and equal. This ensures that the birth and death process is stationary and that the number of active traders is on average constant. Note that, as time is discrete, the values of h_i and τ_i are rounded up to the lowest higher integer.

² All the simulations have been performed using the 'R' computing package, available as freeware at <http://www.r-project.org/> for several platforms. The source code for the benchmark case is available at <http://helios.unive.it/~licalzi/research.html>.

³ The official price is defined as the unweighted mean of all the transaction prices recorded in the last trading session in which the trading volume was positive.

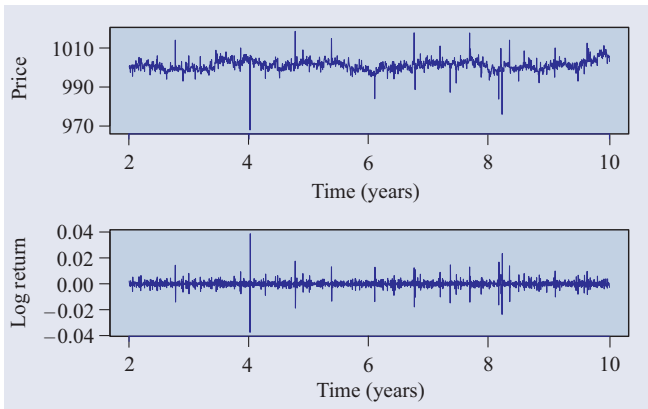


Figure 2. Daily time series for prices (top) and log-returns (bottom).

3.3. Results

The time series for daily official prices and the corresponding log-returns of the benchmark case are presented in figure 2. Transient effects due to the initialization of the simulations usually disappear within 2 months; however, we prefer to be cautious and, throughout this paper, we systematically discard the observations corresponding to the first 2 years (500 trading days). We note that a model based on a minimal set of behavioural assumptions like ours should not be expected to generate ‘realistic’ time series. As is well known in the agent-based literature, this requires a larger set of behavioural assumptions; see for instance Lux (1998). Testing our conjecture requires only to show that no special behavioural assumptions are needed to generate ‘narrow peaks’ and ‘fat tails’ of daily log-returns provided that the structural assumptions are sufficiently detailed.

A close-up of the price time series is given in figure 3. Except for an outlier, prices keep fluctuating approximately within a $\pm 1\%$ band. Note that the price time series tend to exhibit 3 day tents: usually, a decrease in the official price between 2 days is followed by an increase in the following trading session. (Remember that the official price is the unweighted mean of all the transaction prices recorded in 1 day.) This dynamics is consistent with value-based trading, as explained in Farmer and Joshi (2002), and shows up as a significant negative autocorrelation at lag 1 in the daily log-returns.

The viability of the simulated market can be gathered from the top part of figure 4, which plots the histogram of the trading volumes. There are only 48 days in which no trading occurs, corresponding to 2.4% of the trading sessions⁴. The bottom part of figure 4 shows that volumes exhibit a modest but significant positive correlation at lag 1. This is consistent with a market of fundamentalists where, if an agent is rationed, he keeps issuing similar limit orders until he can find a match. Note that we do not restrict agents to trade just one unit per day. This assumption, common to most simulations of order-driven markets in the literature, artificially extends the length

⁴ An increase (respectively, a decrease) of 40% in the number of potentially active traders gave 6 (171) no-trading days, corresponding to 0.3% (8.55%) of the trading sessions.




Figure 3. Daily time series for prices over 1 year.

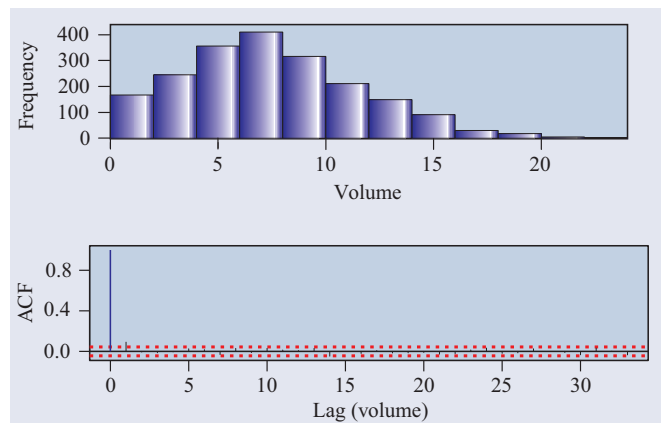


Figure 4. Histogram of trading volumes (top) and volume autocorrelations (bottom), with Bartlett’s confidence bands set at 5%.

of time an agent has to stay rationed and thus generates positive volume autocorrelations at higher lags.

Finally, figure 5 provides a clear test for our conjecture. On the left, we have plotted the histogram of daily log-returns and we have superimposed a best-fit normal distribution. The support has been truncated to about ± 3 standard deviations from the mean to improve readability. The narrow peak is clear-cut and typical of all the simulations we have run. The right part of the figure plots the sample quantiles against the theoretical quantiles of the normal distribution: the divergence for large values reveals fat tails and it is again clear-cut and typical. The Shapiro–Wilk test rejects normality at a confidence level far above $1 - 10^{-8}$ in each of the simulations we have run. The empirical Fisher kurtosis is 43.92.

3.4. Robustness analysis

We have performed an extensive robustness analysis on our model, individually testing each of the parameters in table 1 for sensitivity. We have also specifically examined the probability distributions, replacing them with other unimodal families or modifying their distributional parameters. The narrow peak and the fat tails of daily log-returns always emerge, unless the choice of parameters is so extreme to make the market unviable




Figure 5. Histogram of the distribution of daily log-returns superimposed on the best normal fit (left) and corresponding $Q-Q$ plot (right). The support of the histogram is truncated to improve readability.

for lack of liquidity. This holds even when the specific price dynamics shows long-term effects due to the misalignment of the initial parameters, backing up our claim that a narrow peak and fat tails of short-term log-returns are structural effects.

We present just a few representative examples. (To facilitate comparisons, all the examples presented in this paper use the same seed.) The three time series in the top part of figure 6 illustrate the price dynamics with $c_i = 2200$ (top), $c_i = 2000$ (middle) as in the benchmark case, and $c_i = 1800$ (bottom). Economic theory predicts that the total amount of cash in the system affects positively the overall level of the stock price. We observe that a change of 10% in the cash endowment of each agent accordingly affects the long-term stock price.

Moving to the middle part of figure 6, the time series from top to bottom represent the price dynamics for three different values of the risk premium: $\pi = 0$ (no risk premium); π as in the benchmark case; and $\pi = 0.03$, respectively. These time series have been smoothed to eliminate overlapping spikes and improve readability by applying the variable span smoother described in Friedman (1984). In accord with economic theory, the stock price is higher when the risk premium of agents is (on average) lower.

Consider now the time series in the bottom part of figure 6. The middle time series is the benchmark case. The other two time series are obtained when we offset the initial level of p by $\pm 10\%$. In either case, the system slowly drifts towards its long-term level. The simulations show that our order-driven

market persistently exhibits homeostatic behaviour: the price never settles down to a specific value, but oscillates within a narrow band; when outside of it, the price moves towards this band.

This homeostatic behaviour is illustrated in figure 7. With respect to the benchmark case, we simulated three interventions of monetary policy. We raised the riskless rate from 0.02 to 0.03 at the beginning of the third year, added a further raise from 0.03 to 0.04 at the beginning of year 5, and cut the rate back to its original level at the beginning of year 7. The time series of prices always reacted in the direction predicted by economic theory. More interestingly, once the initial conditions were restored, the price returned to the original long-term level.

Finally, we investigated the dynamics of prices and volumes under a different distribution of agents. Motivated by the pervasive Zipf’s law, we assumed that agents were divided in five classes endowed with different amounts of cash and stock as given in table 2. This particular combination gives a number $n = 3140$ of potentially active traders. This number was used in the simulation of the benchmark case to facilitate a direct comparison. We found no significant differences, but verified that unequal initial endowments tend to increase trading volumes and generate fatter tails for short-term log-returns.




Figure 6. Robustness analysis. Top: benchmark $c_i \pm 10\%$. Middle: $\pi = 0$, benchmark π , $\pi = 0.03$. Bottom: benchmark $p \pm 10\%$.

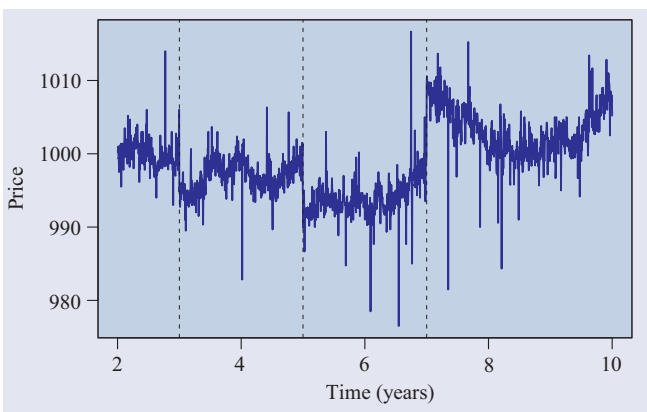


Figure 7. Time series of prices with permanent shocks to the riskless rate of interest: $+0.01$ in year 3, $+0.01$ in year 5 and -0.02 in year 7. (Vertical scale truncated for readability.)

Table 2. Endowments distributed according to a Zipf’s law.

Class	Number of agents	Stock	Cash
1	10×4^4	2^0	2000×2^0
2	10×4^3	2^1	2000×2^1
3	10×4^2	2^2	2000×2^2
4	10×4^1	2^3	2000×2^3
5	10×4^0	2^4	2000×2^4

4. Simulations with zero-intelligence traders

Several of the existing simulations of order-driven markets assume pure noise trading. This is often implemented along the following scheme. A random draw decides whether a trader places a buy or a sell limit order. The order is for a unit quantity of stock, at a price which is usually determined by means of a simple mechanical rule such as ‘increase (or decrease) the current price by one tick’. This trading behaviour provides a rough test for the functioning of an order-driven market, but it forfeits two important restrictions that impair its power to convincingly test our conjecture.

The first restriction is that traders face budget constraints. The second one is that trade should be loosely directed by

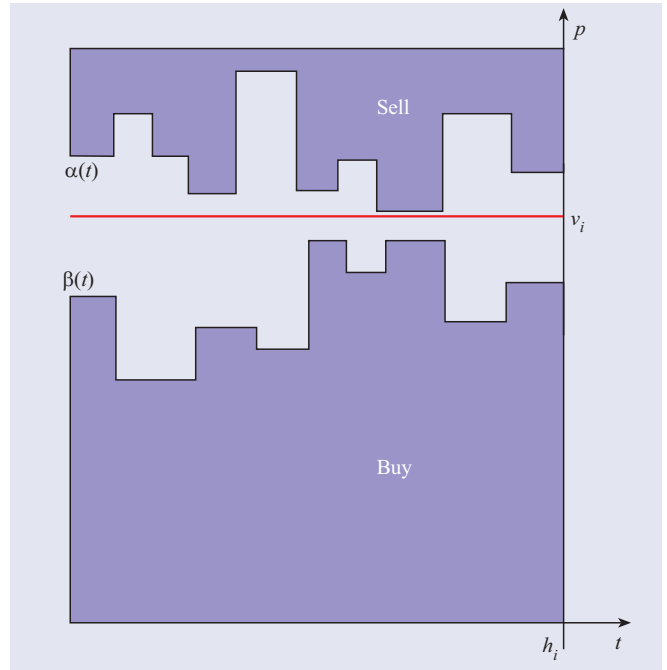


Figure 8. The bid–ask spread.

the purpose of achieving some gain. Our second major test simulates a market with zero-intelligence traders that satisfy this minimal set of behavioural constraints. We maintain all the assumptions and the initializing choices of the benchmark case, but we change the trading rule to the following value-based strategy. When sampled, and compatibly with his budget constraints, a trader places a buy order at a (ticked) price uniformly distributed between $(1 - \Delta)v_i$ and v_i and a sell order at a (ticked) price uniformly distributed between v_i and $(1 + \Delta)v_i$. The bid and ask prices for a trader i are independently renewed at each sample. Hence, trader i repeatedly (and stubbornly) attempts to buy at a price slightly lower than v_i and to sell at a price slightly above v_i . Given the heterogeneity of the v_i values, we believe that this is the minimal set of restrictions that a sensible test for our conjecture should satisfy.

Figure 8 provides a graphical representation of the typical behaviour of a zero-intelligence trader. Given his fundamental value v_i , each day he issues (random) limit orders that identify a (random) band of inactivity around v_i . The trader sells if the price exits the band upward and buys if it exits downward. In terms of the $\beta(t)$ and α functions used for describing the behaviour of fundamentalists, the trader’s bid and ask move up and down in a systematic but unpredictable way.

Figure 9 summarizes this second test of our conjecture. The top part reports the time series of prices for $\Delta = 5\%$. On the left side of the bottom part, we have plotted the histogram of daily log-returns and we have superimposed a best-fit normal distribution. The support has been truncated to about ± 3 standard deviations from the mean to improve readability. The narrow peak shows up again and it is typical of all the simulations we have run with zero-intelligence traders. On the right side, we have plotted the sample quantiles against the theoretical quantiles of the normal distribution: the divergence

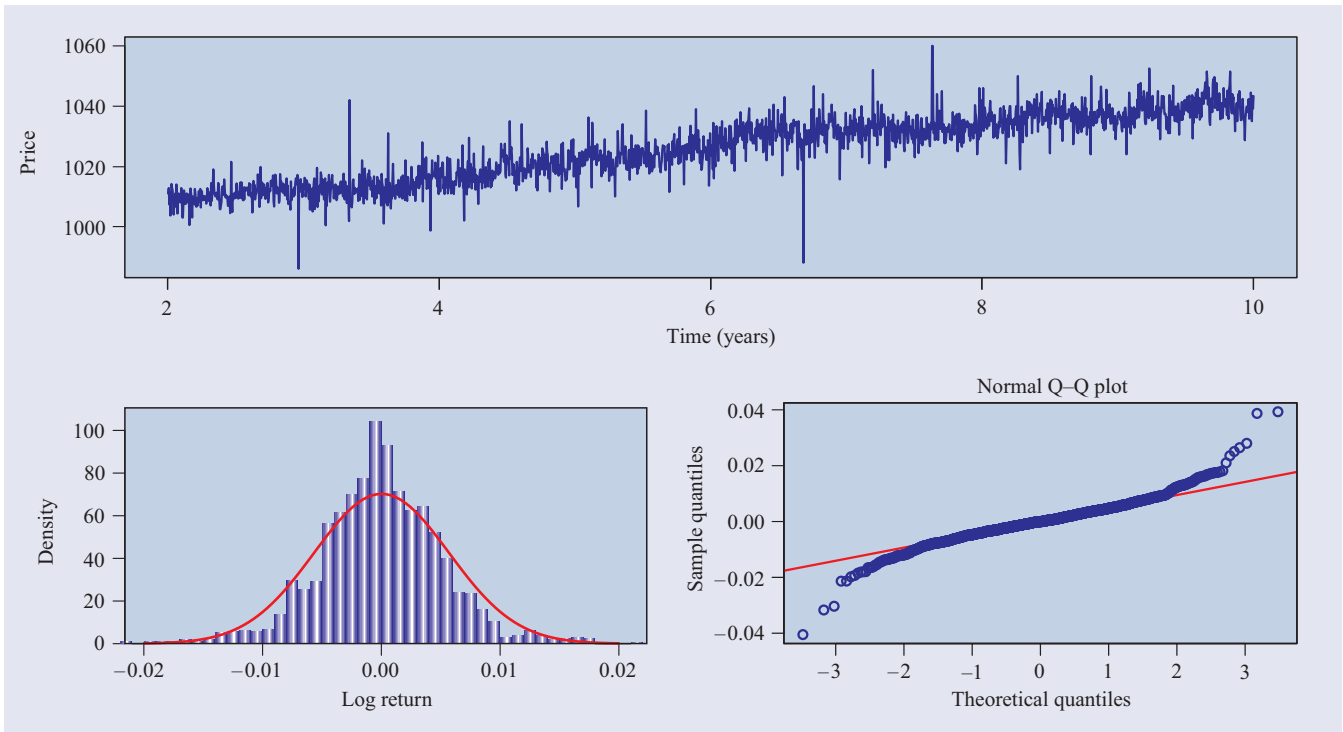


Figure 9. Zero-intelligence traders: price time series (top), histogram of the daily log-returns (bottom, on the left) and $Q-Q$ plot (bottom, on the right).

for large values reveals fat tails and it is again clear-cut and typical. The Shapiro–Wilk normality test rejects normality at a confidence level far above $1 - 10^{-8}$ in each of the other simulations we have run. The empirical Fisher kurtosis is 8.44.

The robustness analysis for the case of zero-intelligence traders is simpler than the analysis for fundamentalist traders, because there are a few parameters such as r and π that are now irrelevant and can be ignored. Except for these obvious cases, however, we performed the same long series of individual testing of each parameter and each distributional assumptions. Invariably, a narrow peak and fat tails of daily log-returns have emerged throughout.

It is instructive to contrast the price time series for the benchmark case in figure 2 with the time series on the top part of figure 9. The only difference between the two cases is the behavioural rule followed by traders. Yet, while the first time series fluctuates around a roughly constant level, the second one has a clear upward drift. This difference in behaviour is easily explained by comparing figure 1 with 8. For equal v_i values, fundamentalists tend to place bids and asks at a lower price than zero-intelligence traders. Hence, the stationary price level in a market populated with the latter strategy is higher and figure 9 shows the system chasing its own stationary level.

More importantly, this comparison underlines that the fundamentalists of section 3 and the zero-intelligence traders of section 4 share two important restrictions on behaviour: they both apply value-based trading and they both satisfy the budget constraints. The difference between the two classes of traders concerns only the rules governing their bid–ask spread: zero-intelligence trading requires no understanding of the fundamentals of the economy (except for the estimate

v_i); fundamentalist trading requires processing other pieces of information such as the time-value of money or the risk attitude, which are here r and π_i , respectively.

The limiting case for both trading rules leads to the ‘plain vanilla’ modelling assumption for fundamentalist behaviour: buy if the current price is strictly below v_i and sell if it is strictly above. This obtains as the limit of our fundamentalists’ bid and ask functions (1) and (2) for $r \downarrow 0$ and $\pi \downarrow 0$, or as the limit of the zero-intelligence rule for $\Delta \downarrow 0$. In this respect, one can view our model as a noisy version of a system with plain vanilla fundamentalists. The introduction of a degree of heterogeneity in the traders’ parameters leads to an increase in the noise of the system. However, it should be stressed that neither our fundamentalist nor our zero-intelligence traders are noise traders (such as those in Maslov (2000), for instance) because they are forced to apply value-based trading and to satisfy liquidity constraint. In other words, we have noise in the system but the agents are not noise traders.

Based on this, one might wonder whether our claim that a structurally detailed model of an order-driven market suffices to generate fat tails and narrow peaks would still hold in an environment populated only with plain vanilla fundamentalists. To answer this legitimate concern⁵, we tested the market assuming that fundamentalists follow the simplest possible trading rule: buy (subject to budget constraints) as much as possible whenever the current price is strictly below v_i and sell (subject to budget constraints) as much as possible whenever it is strictly above. The simulation presented in figure 10 uses the same parameters of the benchmark case

⁵ We thank the referees for raising it.

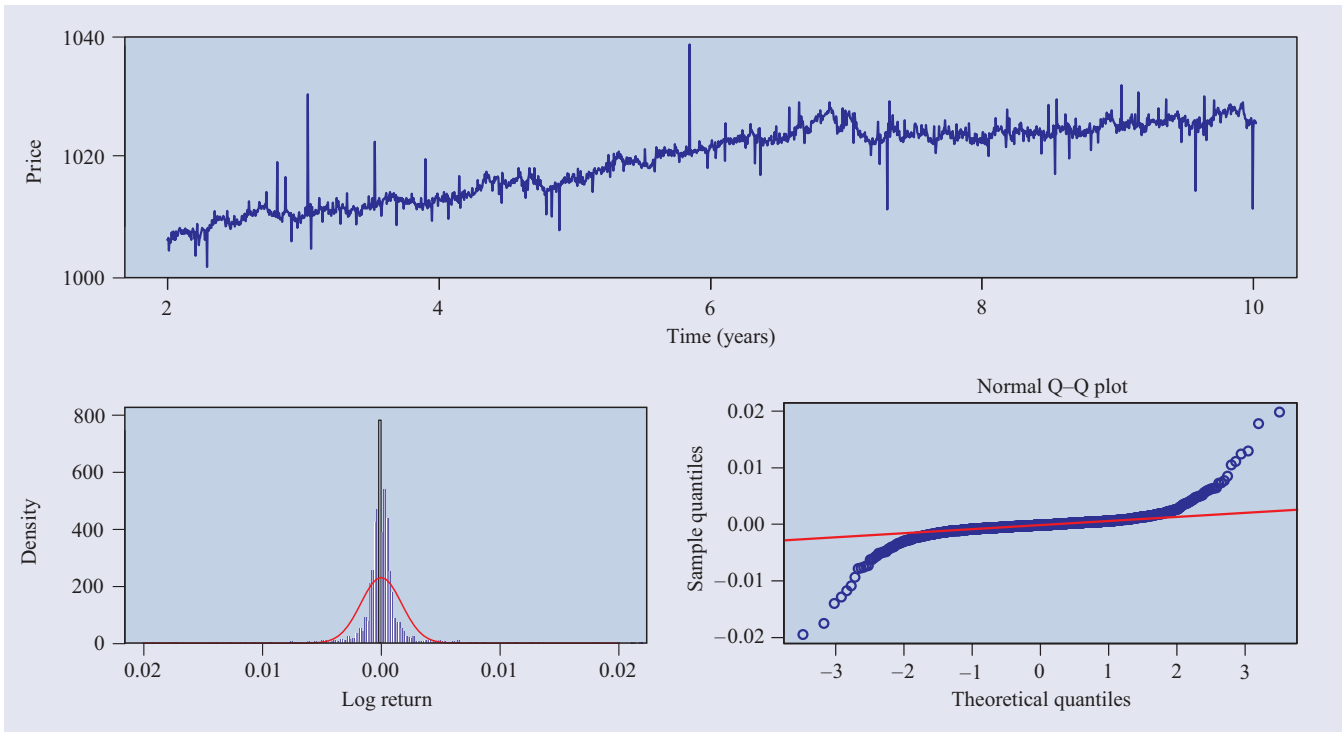


Figure 10. Plain vanilla fundamentalists: price time series (top), histogram of the daily log-returns (bottom, on the left) and $Q-Q$ plot (bottom, on the right).

except for setting $r = \pi = 0$ throughout. The emergence of fat tails and narrow peaks is once again clearly confirmed.

5. Conclusions

We have modelled the precise functioning of an order-driven market and simulated its behaviour under two simple but realistic behavioural rules, respectively associated with fundamentalist and zero-intelligence value-based trading under budget constraints. Unless the market is unviable for lack of trading volume, all simulations consistently give a leptokurtic distribution of daily log-returns which exhibits both a narrow peak and fat tails. These two properties are so pervasive that we advance the claim that they are a consequence of the structural properties of the market rather than of the fine-grained details of the agents' behavioural rules.

This claim should not be used to downplay the importance of behavioural assumptions in achieving a realistic description of financial time series. There are many other well known properties, such as volatility clustering or short-term correlations, which the market structure alone simply cannot explain. The agent-based literature has already uncovered several behavioural assumptions that help to achieve time series that are more realistic than ours. One of the next steps will probably be to study how behavioural assumptions and market structure combine to yield the characteristic 'signature' of different financial markets.

Acknowledgments

We thank three anonymous referees and participants at CEF2002 for their comments. Financial support from MIUR is acknowledged.

References

- Arthur W B, Holland J H, LeBaron B, Palmer R and Tayler P 1997 Asset pricing under endogenous expectations in an artificial stock market *The Economy as an Evolving Complex System II* ed W B Arthur, S N Durlauf and D A Lane (Reading, MA: Perseus Books) pp 15–44
- Bak P, Paczuski M and Shubik M 1997 Price variations in a stock market with many agents *Physica A* **246** 430–53
- Beja A and Goldman M B 1980 On the dynamic behaviour of prices in disequilibrium *J. Finance* **35** 235–48
- Bouchaud J-P, Mézard M and Potters M 2002 Statistical properties of stock order books: empirical results and models *Quant. Finance* **2** 251–6
- Chiarella C and Iori G 2002 A simulation analysis of the microstructure of double auction markets *Quant. Finance* **2** 346–53
- Cohen K J, Maier S F, Schwartz R A and Whitcomb D K 1978 Limit orders, market structure, and the returns generation process *J. Finance* **33** 723–36
- Daniels M G, Farmer J D, Gillemot L, Iori G and Smith E 2002 A quantitative model of trading and price formation in financial markets *Preprint cond-mat/0112422* December
- Day R H and Huang W 1990 Bulls, bears and market sheep *J. Econ. Behav. Organ.* **14** 299–329
- Farmer J D and Joshi S 2002 The price dynamics of common trading strategies *J. Econ. Behav. Organ.* **49** 149–71
- Friedman J H 1984 A variable span smoother *Technical Report 5* Laboratory for Computational Statistics, Department of Statistics, Stanford University

Q.1

- Gode D K and Sunder S 1993 Allocative efficiency of markets with zero intelligence traders *J. Political Economy* **101** 119–37
- Lux T 1998 The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distribution *J. Econ. Behav. Organ.* **33** 143–65
- Lux T and Marchesi M 2000 Volatility clustering in financial markets: a micro-simulation of interacting agents *Int. J. Theor. Appl. Finance* **3** 675–702
- Maslov S 2000 Simple model of a limit order-driven market *Physica A* **278** 571–8
- Matassini L and Franci F 2001 On financial markets trading *Physica A* **289** 526–42
- Mendelson H 1982 Market behaviour in a clearing house *Econometrica* **50** 1501–24
- Milgrom P and Stokey N 1982 Information, trade and common knowledge *J. Econ. Theory* **26** 17–27
- Raberto M, Cincotti S, Focardi S M and Marchesi M 2001 Agent-based simulation of a financial market *Physica A* **299** 319–27
- Smith E, Farmer J D, Gillemot L and Krishnamurthy S 2002 Statistical theory of the continuous double auction *Preprint cond-mat/0210475* October
- Steiglitz K and Shapiro D 1998 Simulating the madness of crowds: price bubbles in an auction-mediated robot market *Comput. Economics* **12** 35–59
- Stigler G J 1964 Public regulation of the securities markets *J. Business* **37** 117–42
- Welch I 2000 Herding among security analysts *J. Financial Economics* **58** 369–96