Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Physica A 359 (2006) 619–633
PHYSICA A
www.elsevier.com/locate/physa

Breeds of risk-adjusted fundamentalist strategies in an order-driven market

Marco LiCalzi, Paolo Pellizzari*

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

Received 12 November 2004; received in revised form 9 June 2005
Available online 13 July 2005

Abstract

This paper studies an order-driven stock market where agents have heterogeneous estimates of the fundamental value of the risky asset. The agents are budget-constrained and follow a value-based trading strategy which buys or sells depending on whether the price of the asset is below or above its risk-adjusted fundamental value. This environment generates returns that are remarkably leptokurtic and fat-tailed. By extending the study over a grid of different parameters for the fundamentalist trading strategy, we exhibit the existence of monotone relationships between the bid–ask spread demanded by the agents and several statistics of the returns. We conjecture that this effect, coupled with positive dependence of the risk premium on the volatility, generates positive feedbacks that might explain volatility bursts.

Keywords: Price dynamics; Statistical properties of returns; Market microstructure; Agent-based simulations

1. Introduction

The agent-based paradigm has sparked a wealth of models that try to explain many well known stylized facts of price dynamics in financial markets. These models use rich sets of hypotheses about the behavioral strategies of the agents and the

*Corresponding author. Tel.: +39 041 234 6924; fax: +39 041 522 1756.
E-mail addresses: licalzi@unive.it (M. LiCalzi), paolop@unive.it (P. Pellizzari).

0378-4371/$ - see front matter
doi:10.1016/j.physa.2005.06.039

market microstructure underlying the mechanism of price formation. Among the former, there are several known variants of value-based fundamentalist, noisy, and chartist trading strategies, either with or without explicit optimization of an objective function, compounded by communication, imitation and synchronization effects. Among the latter, the most important variations concern the imposition of budget constraints or the fine details of the trading protocol, ranging from tâtonnement adjustments to the rarer but more realistic case of order-driven markets.

Many models in the literature are able to reproduce some (but seldom all) the reported stylized facts, such as leptokurtic and fat-tailed distribution of returns, volatility effects and volume persistence. However, a unifying view of agent-based markets has still to come and it is difficult to understand unambiguously which ingredients of a given model are the driving cause of the empirically observed features.

This paper builds on [1], where it is shown that the interaction of purely fundamentalist agents with heterogeneous estimates in a order-driven market is sufficient to generate fat tails and a leptokurtic shape of the return distribution. We confirm the robustness of this conclusion under an alternative formulation and across a wide range of parameter values. The main contribution is to show that the bid–ask spread demanded by the agents has a key role in shaping some macroscopically observable properties of the returns' time series, such as its excess kurtosis, volatility and tail exponent. We find that there are monotone relationships between the agents' bid–ask spreads and these important statistics.

This monotonicity suggests an alternative explanation for volatility bursts that is not based on the characteristics of the information flow or coupling effects (such as imitation) among agents. The bid–ask spread demanded by an agent is directly affected by his risk aversion and by his propensity to a higher efficacy (versus a higher immediacy). If agents raise their risk premia or their efficacy requirements in reaction to an increase in volatility, they widen their bid–ask spreads. As a consequence, if the trading environment generates a volatility which is increasing in the bid–ask spreads, we have a positive-feedback mechanism that can generate volatility bursts without further assumptions. We strive to keep our model as simple as possible precisely to help disentangling how much of this effect is attributable to our behavioral or structural assumptions.

Our work relates to the literature in many important ways. Order-driven (also known as book-based) markets have been studied in Ref. [2–4]. There is a mounting evidence that a detailed modeling of the market mechanism is important to understand the statistical properties of the returns' time series; see Ref. [5–7]. The model developed in Ref. [3] shows that noisy trading, generated by random orders anchored to the current price, can produce fat-tailed returns. Our model shows that the same conclusion holds when more realistic agents are used; moreover, it removes the restriction (present in other papers as well) that an agent can trade at most one unit of stock at a time.

The rich market model in Ref. [2] includes both noisy and value-based traders interacting in a order-driven market in which both imitative effects and volatility feedbacks are explicitly considered. Fundamentalists buy (or sell) stocks when the

expected dividend rate exceeds (or is smaller than) the interest rate by a margin Δ which is called "stickiness". While the results of the model are complex and diverse, they are the outcome of so many interacting variables that it is difficult to ascertain which ones are more important in determining the overall outcome. This problem is exacerbated, for instance, in Ref. [4], where budget-constraints and book trading are coupled with *networks of friends* and an embarrassing number of behavioral parameters (personal target price, expected gain, maximum loss, personal stop-loss, …).

A different strand of papers eschews order-driven markets in favor of alternative price formation mechanisms, based on matching supply and demand [8], or Walrasian auctioneer approximations [8–11]. These papers model herd effects or other related synchronizations in trading strategies using random graphs or lattices, which indeed produce fat tails and volatility bursts in the returns' time series. They leave it an open issue whether fat tails can be obtained in the absence of any kind of herding phenomenon. By contrast, the model presented in this paper does not allow agent to share knowledge or to coordinate in any way: the only common information used by the agents is the current price. Based on this, we do not expect to see heteroskedasticity in the returns' time series: here, our result is weaker than what others have obtained; on the other hand, we are able to rule out communication or other informational effects as the source of the fat tails generated by our model.

Our fundamentalist agents buy (or sell) stock whenever it is under- (or over-) valued with respect to their own individual risk-adjusted estimate of the fair value. Besides budget constraints, the determinants of the trading decisions for all agents are just four: their estimate of the fundamental value of the risky asset, their investment horizon, the risk premium they require per unit of time, and their propensity to search efficacious trading. Agents are heterogeneous with respect to these four parameters. On the other hand, they are oblivious to external news or to the trading patterns of other agents. These simple-minded or perhaps "stubborn" fundamentalist agents, when interacting in an order-driven market, suffice to robustly generate leptokurtic and fat-tailed returns. This provides a useful benchmark in the search for the minimal set of assumptions ensuring the emergence of these well known statistical properties.

Our study of different breeds of agents who share the same trading rule but have different bid–ask spreads has uncovered the existence of monotone relationships between the (average) demanded spread on one side and excess kurtosis, tail exponents and volatility of the returns on the other side. An increase in the demanded spread in the traders' population appears to cause an increase in the volatility of the returns and a decrease in the excess kurtosis and in the tail exponent. We conjecture that these relationships can be used to explain other effects, when appropriate feedback mechanisms are taken in due consideration. For instance, suppose that agents form their preferences in a way that increases their demanded spreads in reaction to (possibly local) increases in the volatility of returns. Then, small noisy upward deviations in volatility would translate in higher spreads that could in turn spark a further increase in volatility, starting an escalation analogous to a volatility burst.

The paper is organized as follows. Section 2 describes the agents and the market model in a compact way; see [1] for additional details. Section 3 presents our simulations and discuss the results. Finally, Section 4 offers some concluding remarks.

2. The model

We consider an economy with two assets. One is a bond that pays a riskless yearly interest rate r. The other is a stock that has a risky price p and pays no dividends. No new information is ever released. 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. Agents are budget-constrained: they are not allowed to short sell or to borrow money.

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 on the order of 1/1000 of the stock price.

The value-based trading strategy used by the agents is the following. At any point in time, each agent has two thresholds: when the price of stock is below the lower threshold, he wants to load on stock (compatibly with budget constraints); when this price is above the upper threshold, he wants to load on bond (compatibly with budget constraints); when the price of stock is between the two thresholds, he stays put. The lower threshold represents the highest bid price $\beta_i(t)$ that an agent is willing to offer at time t and the upper threshold is the lowest ask price $\alpha_i(t)$ that he is willing to accept at time t. Clearly, the bid (ask) price of an agent bounds from above (below) the best purchase (sell) limit order that an agent would place.

We operationalize the formation of bids and asks as follows. Each agent formulates his own assessment v_i of the fundamental value of the stock. Each agent is a fundamentalist in the sense that, once v_i has been estimated, it does not change anymore. In particular, v_i is not affected by the dynamics of the stock price p and thus it can be thought of as a constant trait of agent i. 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 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 load on stock at time t if he expects a return sufficiently

higher than the return on the riskless bond; that is, if

$$\frac{v_i}{p} \geqslant \mathrm{e}^{(r+\pi_i)\cdot(h_i-t)} .$$

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

$$\beta_i(t) = v_i \mathrm{e}^{-(r+\pi_i)\cdot(h_i-t)} . \tag{1}$$

Similarly, trader i prefers to load on bond when the risk-adjusted expected return from holding stock is lower than the riskless return achievable from investing in the bond. The risk premium is $\pi' < \pi$, because the agent is moving away from risk. Hence, the lowest ask price he is willing to accept at time t is

$$\alpha_i(t) = v_i \mathrm{e}^{-(r+\pi'_i)\cdot(h_i-t)} . \tag{2}$$

To cut down on the number of free parameters, we parsimoniously assume $\pi' = \pi/2$. Except for the explicit time-dependence, this behavioral rule is similar to the one suggested in Ref. [2] for the fundamentalists' trading decisions; Ref. [12] provides a strong case for the use of thresholds in modeling agent-based financial markets.

The spread between bid and ask may be affected by another important component, which is known as the attitude towards the immediacy versus efficacy tradeoff. Consider for instance $\beta_i(t)$. Agent i would never submit a purchase limit order at a price higher than $\beta_i(t)$. However, in the attempt to fetch a more favorable transaction price, the agent may choose to offer a lower bid (or perhaps to wait for more favorable price conditions). Doing so, of course, reduces the chances to complete a transaction soon (decreasing immediacy) but increases its profitability (improving efficacy). In general, we should not expect an agent to place a purchase limit order at his "bid" price $\beta_i(t)$ but at some lower price $[1-\delta_i(t)]\beta_i(t)$ where $\delta_i(t)$ in $[0,1]$ is some appropriate "shading" factor. Similarly, a sell limit order should be made at some higher price $[1+\delta'_i(t)]\alpha_i(t)$, for some $\delta'_i(t) \geqslant 0$. In the agent-based literature, this point was made clear in Ref. [13], where zero-intelligence agents choose randomly by which shading factor to adjust their valuations before attempting a trade; see also Ref. [14], where agents try to get a "profit level" in excess of their limit prices. To cut down on the number of free parameters, we parsimoniously assume $\delta'_i(t) = \delta_i(t)$.

When $t = h_i$, the agent has reached his investment horizon and exits the market. His endowment 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. Investment horizons, risk premia and shading factors are randomly distributed across traders. Horizons and premia are viewed as constant traits of an agent, which are chosen when he becomes active and remain constant until he exits the market. Shading factors are resampled every time the agent attempts a transaction.

The theoretical bid–ask spread $\alpha_i(t) - \beta_i(t)$ is smaller than the actual bid–ask spread $[1+\delta_i(t)]\alpha_i(t) - [1-\delta_i(t)]\beta_i(t)$. In the following, we use bid–ask spread to refer exclusively to the actual bid–ask spread demanded by an agent. This bid–ask spread is decreasing in the residual time $(h_i - t)$ and it is increasing in π_i and $\delta_i(t)$.

Therefore, an agent is likely to trade more actively towards the end of his activity period; moreover, increasing the general level of agents' risk premia or bid–ask spreads tends to increase the volatility of the transaction prices by making them bounce in a larger corridor.

Trading over the book takes place as usual. Let a be the best ask price available when trader i checks the book. If $(1-\delta)\beta_i(t) \geqslant a$, 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 buys only 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 nth level ask price on the book lower than $(1-\delta)\beta_i(t)$. If this second condition occurs, the agent places a limit order at a price equal to the first tick below $(1-\delta)\beta_i(t)$ for the maximum quantity of stock he can afford to pay for. Trading in the other direction takes place similarly. We assume that each trading session lasts one day. Since there are about 250 trading days in a year, we clock time in increments of $\frac{1}{250}$. Within one day, all traders check the book and place orders asynchronously. We randomize with uniform probability the order in which traders check the book within one day and we clear the book at the end of a trading session. The book matches orders using price priority and, in case of equal prices, temporal priority.

3. Simulations

This section reports on the simulation of various markets populated by the fundamentalist agents described in the previous section. We view our procedure as the simulation of a (discretely sampled) continuum of markets, which are parsimoniously indexed by two parameters that shape the distribution of bid–ask spreads demanded by the agents.

Each time a new agent enters the market, we randomly draw his individual risk premium π_i according to the uniform distribution on the interval $[0, \Pi]$, where $\Pi \geqslant 0$ is a global parameter identical across agents. While π_i is resampled (once and for all) for each new agent, the overall impact of the risk premia can be appraised by Π, which yields an individual average risk premium of the order of $\Pi/2$. Similarly, each time an agent attempts a transaction, we randomly draw his shading factor $\delta_i(t)$ according to the uniform distribution on the interval $[0, \Delta]$, where $\Delta \geqslant 0$ is a global parameter identical across all agents. While $\delta_i(t)$ is resampled every time, the overall impact of the shading factors can be evaluated by Δ.

The two global constants Π and Δ define a continuum of markets $M(\Pi, \Delta)$. In this respect, the analysis reported below can be interpreted as a test of robustness which shows that the conclusions in Ref. [1] hold in a variety of situations and are not an artifact of specific parametric values. Alternatively, the agents trading in our simulations can be thought as belonging to different two-parameter breeds of the same behavioral family $B(\Pi, \Delta)$. A similar approach was used in Ref. [15], where the vector (Π, Δ) is the genotype of the whole family of traders. Under this viewpoint, perhaps more interesting, the study of different markets offers useful insights to

understand how risk premia and shading factors affect the macroscopic properties of the time series of returns.

Table 1 summarizes the parameters of the model. 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; all the individual parameters are reset each time a trader is activated, except for $\delta_i(t)$ which is resampled every time the trader attempts a transaction. The parameter τ_i determines how many periods it takes before a new agent replaces an old agent who has exited the market.

This choice of parameters is loosely inspired by the empirical evidence, but we do not attempt to calibrate the model to a specific set of real data. Our simulations are meant to show that a simple order-driven market can generate both fat tails and a narrow peak in the distribution of daily returns even if all the traders in the market are fundamentalists. Here and in the rest of the paper, we follow custom and define the return on stock as $\rho_t = \log(p_t/p_{t-1})$. The price p_t is measured as the simple (i.e., not weighted by volume) average of the prices at which trade has occurred within a trading session t. If no transaction occurred during the session, p_t is set equal to p_{t-1}.

We are interested in understanding how the determinants of the bid–ask spread demanded by each agent impact on the statistical properties of the returns and on other relevant aspects of the market, such as traded volume or global indicators for its riskiness.

We have simulated 5500 trading days for each combination of values for Π and Δ. Although transient initialization effects disappear very quickly (less than 20 sessions), we have followed a precautionary principle and we systematically report statistics after discarding the first 500 returns. Therefore, all the subsequent analyses are based on time series of 5000 returns, roughly equivalent to 20 years of daily data. To help visual comparisons, we have used the same random seed in every simulation.

Table 1
Identification of the model

	Parameters		Initialization	Label
Global	n	$=$	3410	Number of agents
	r	$=$	0.01	Riskless interest rate
	Π	$\in$	$\{0, 0.03, 0.06, 0.09, 0.12, 0.15\}$	Risk premium
	Δ	$\in$	$\{0, 0.025, 0.05, 0.075, 0.1\}$	Shading factor
Trader	c_i	$=$	2000	Initial cash endowment
	s_i	$=$	1	Initial stock endowment
	v_i	$\sim$	$1000 \cdot [1 + \mathrm{U}(-0.1, 0.1)]$	Fundamental value
	π_i	$\sim$	$\mathrm{U}(0, \Pi)$	Risk premium
	$\delta_i(t)$	$\sim$	$\mathrm{U}(0, \Delta)$	Shading factor
	h_i	$\sim$	$t + \lceil \mathrm{Exp}(1/250) \rceil$ days	Investment horizon
	τ_i	$\sim$	$h_i + \lceil \mathrm{Exp}(1/250) \rceil$ days	Refresh time

Table 2
Summary of results

1. Even in an order-driven market populated only by fundamentalists, leptokurtosis and fat tails in the empirical distribution of returns are robust features.
2. The excess kurtosis and the tail exponents of the distribution of returns are decreasing in the bid–ask spread demanded by the agents.
3. The volatility of returns is increasing in the bid–ask spread demanded by the agents.
4. The downside risk is increasing in the bid–ask spread demanded by the agents.
5. The traded volume is decreasing in the bid–ask spread demanded by the agents.

A concise summary for the results of the simulations we have performed is reported in Table 2. The first finding confirms the robustness of a result already exhibited in Ref. [1]. The other four concern the comparative statics of a few summary statistics with respect to the bid–ask spread demanded by the agents, which we proxy by Π and Δ. In particular, we assess the downside risk using two customary measures such as the Value At Risk (VaR) and the Expected Shortfall (ES) for the series $\hat{\rho}_t$ of the simulated returns. The VaR at a confidence level α is the value x such that $Freq(\hat{\rho}_t \leqslant -x) = \alpha$; given a VaR, the ES is $E[\hat{\rho}_t | \hat{\rho}_t \leqslant -VaR]$.

It is worthwhile noting that these findings imply that an increase in the determinants of the bid–ask spread demanded by the agents inflates the volatility of the returns and the downside riskiness of the market, but at the same time tends to reduce the leptokurtosis of returns.

The rest of this section reports on the evidence backing the results summarized in Table 2. We begin with Fact 1. Fig. 1 represents the typical time series, density and normal quantile plot of the returns generated for $\Pi = 0.12, \Delta = 0$ and $\Pi = 0, \Delta = 0.025$, respectively. The first time series has parameters inspired from the equity premium puzzle first described in Ref. [16], that reports a historical risk premium for NYSE stocks of around 6%. As our agents independently and uniformly draw their π in $[0, 0.12]$, the average risk premium is very close to this value. These two time series have nearly the same standard deviation (0.0036 and 0.0040). They are also remarkably leptokurtic and extremely fat tailed, as shown by the normal quantile plot. These general characteristics continue to hold over a wide range of values for Δ and Π.

Consider now Fact 2. Fig. 2 shows the surface of the (empirical[1]) excess kurtosis for all different combinations of Π and Δ. (The z-axis is truncated to improve readability.) The (empirical) excess kurtosis is always positive: our estimates spanned the interval $[1.91, 101.38]$, with mean 19.14 and median 5.01. The graph suggests that an increase in either Π or Δ results in a smaller (empirical) excess kurtosis. Similar relationships hold for the tail exponents of the empirical distribution of returns. When we approximate the tails of the empirical distribution by $x^{-(\mu+1)}$, we get estimates of the magnitude μ of the tail exponent in the range $[2.03, 6.76]$. Fig. 3 depicts the surface of μ for all different combinations of Δ and Π. It is apparent that tails get fatter in correspondence of lower values of Δ and Π. Fig. 4 shows the scaling

[1]Our estimates of the tail exponent imply that the theoretical kurtosis might not be finite.

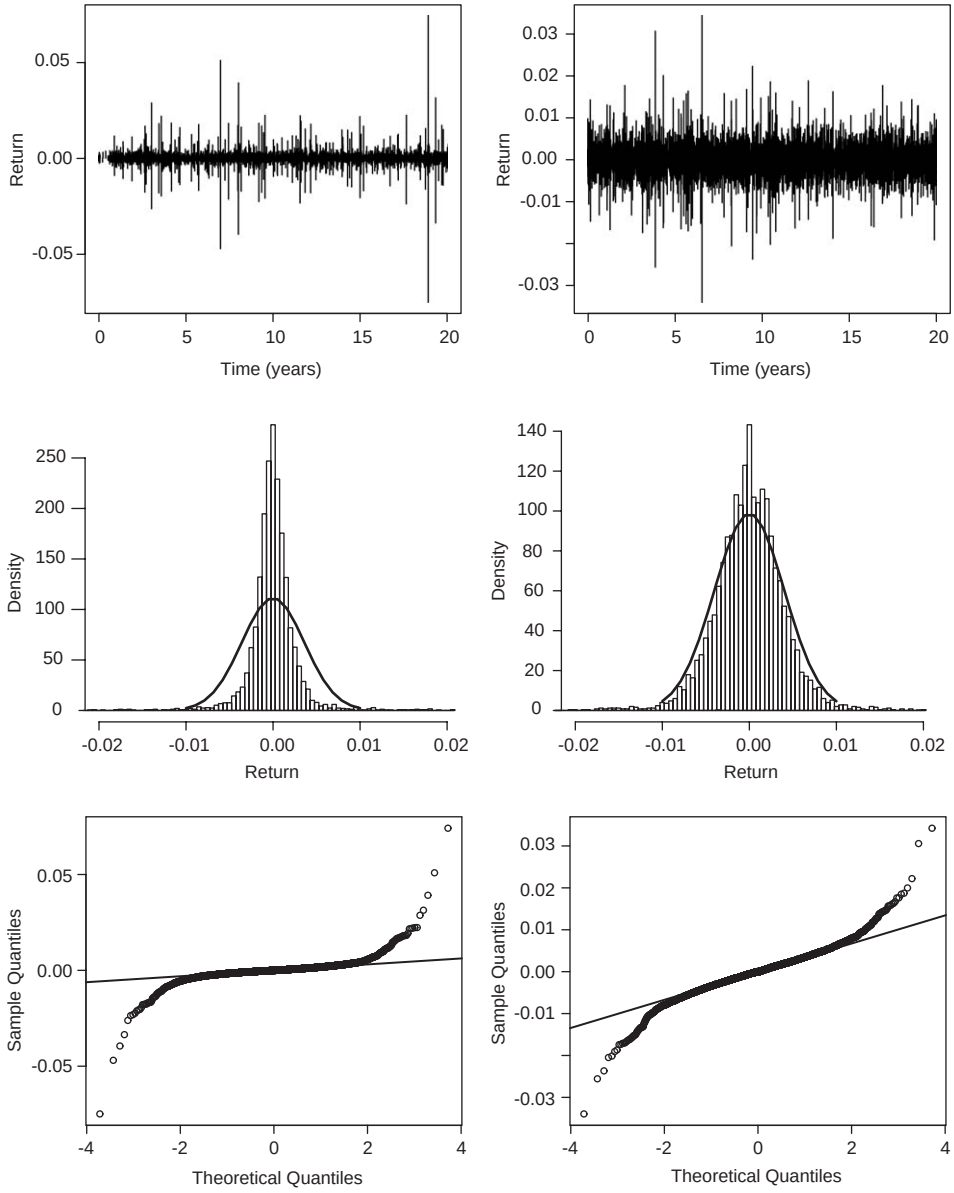

Fig. 1. Time series, density and normal quantile plot for the return time series, when $\Delta = 0, \Pi = 0.12$ (left) and $\Delta = 0.025, \Pi = 0$ (right).

behavior of the tail exponent for the "Mehra-Prescott" case where $\Delta = 0$ and $\Pi = 0.12$. The x-axis holds the returns scaled by their standard deviation and the y-axis reports the probability that the absolute return exceeds a given value. The power

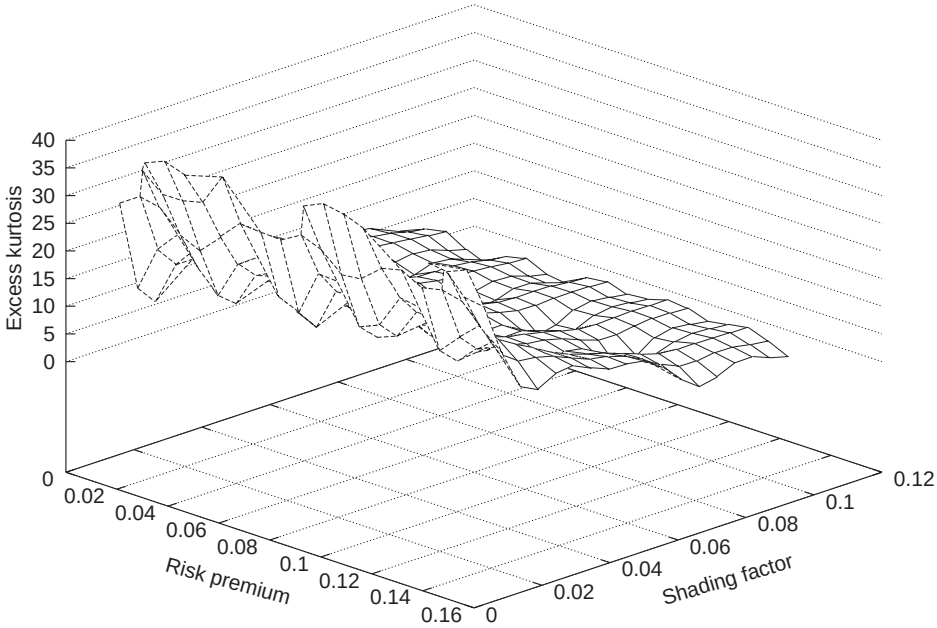

Fig. 2. Kurtosis of the return time series for different combinations of Δ and Π.

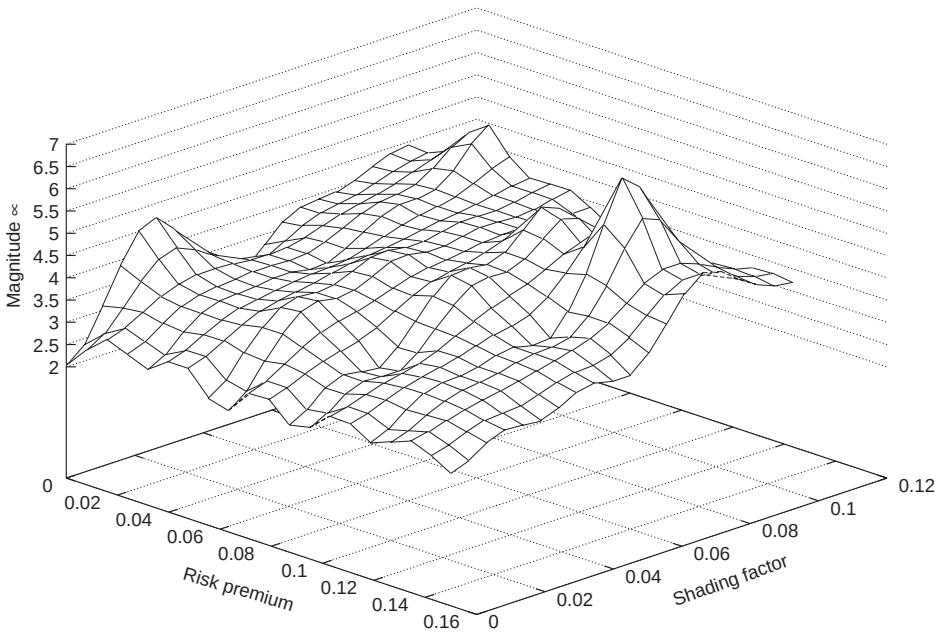

Fig. 3. The magnitude μ for different combinations of Δ and Π.

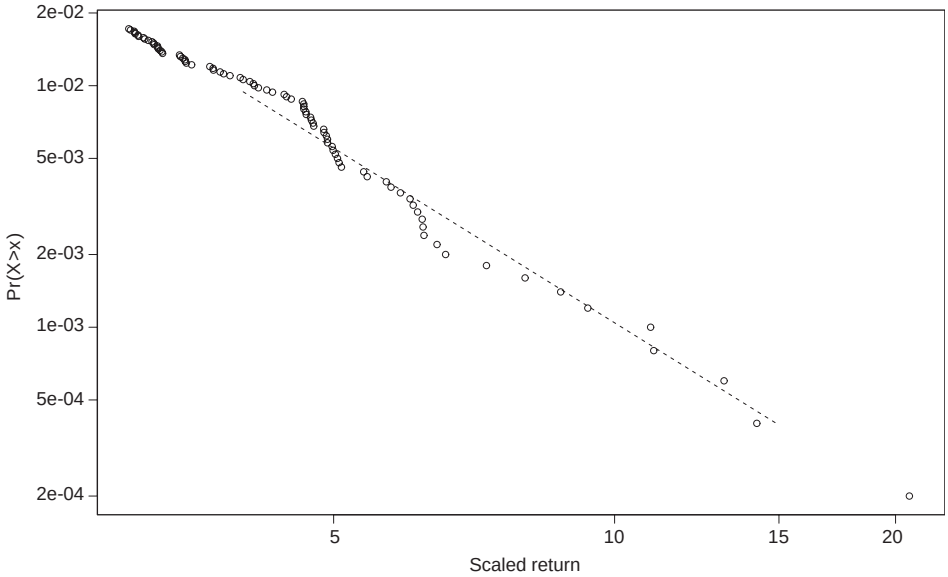

Fig. 4. Scaling properties of the simulated returns when $\Delta = 0, \Pi = 0.12$.

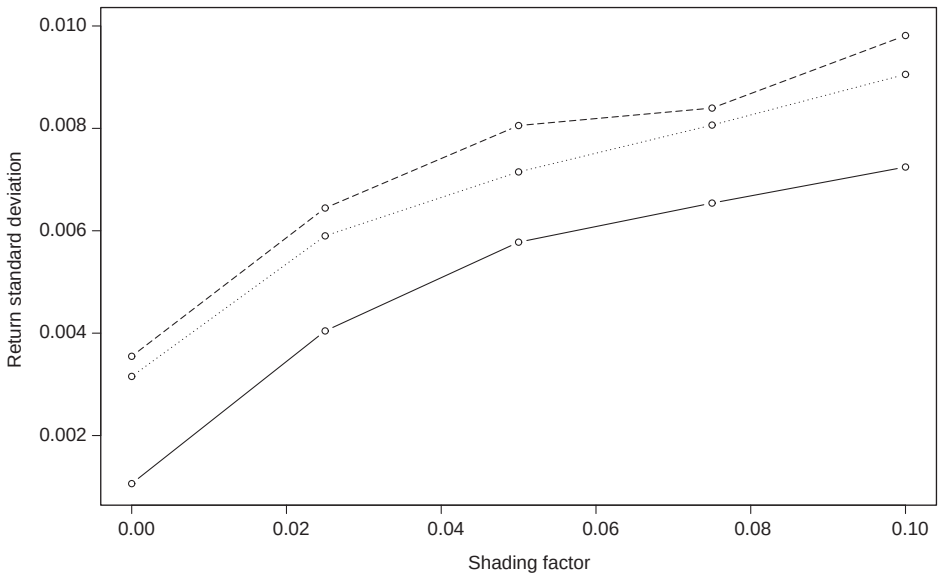

Fig. 5. Standard deviation of the return time series for different values of Δ and Π (from bottom to top, $\Pi = 0, 0.06, 0.12$).

law exponent is the slope of the straight line in the figure. Its estimated value is $\hat{\mu} = 2.40$ and falls within the usual range $[2, 4]$ generally observed in empirical data.

Consider now Fact 3. Fig. 5 exhibits the standard deviation of the returns as a function of Δ, for different values of Π. The volatility of the returns (measured by the standard deviation) is increasing in Δ and Π. The combination of Facts 2 and 3 suggests that the bid–ask spread demanded by the agents impacts differently on the volatility and on the "normality" of returns. While a larger spread tends to increase price bounces between bid and ask and therefore raises volatility, it also makes systematic departures from normality less marked. There is more variability, but less large deviations.

This makes the overall evaluation of the riskiness of the market less clearcut. On the other hand, the literature on risk measurement has established that a more informative evaluation of risk should especially insist on the downside risk. The two best known measures are the VaR and the ES. While the VaR is simpler to compute and therefore far more popular among practitioners, the ES is generally regarded as a better risk measure because it satisfies coherence; see Ref. [17]. We have computed both indicators, obtaining the same qualitative results. For the sake of brevity, we report here only the results for the ES indicator. Fig. 6 graphs the ES for the returns at the 0.01 confidence level. Consistently with our Fact 4, the ES is increasing with Δ and Π. This matches Fact 3, which shows that the standard deviations associated to (log)-returns are also increasing in Δ and Π. Given that the empirical distributions in Fig. 1 are nearly symmetric, this is not surprising. Taken together, Facts 3 and 4 confirm that the riskiness of the market increases with the bid–ask spread demanded by the agents. This conclusion is robust to using alternative measure of riskiness such standard deviation, VaR, or ES.

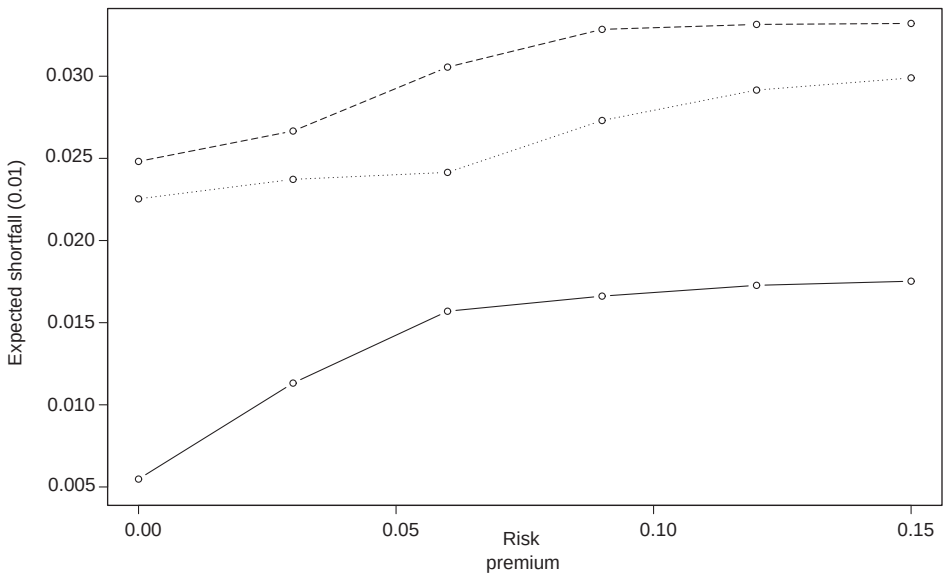

Fig. 6. Expected shortfall at 0.01 confidence level for different values of Δ and Π (from bottom to top, $\Delta = 0, 0.05, 0.1$).

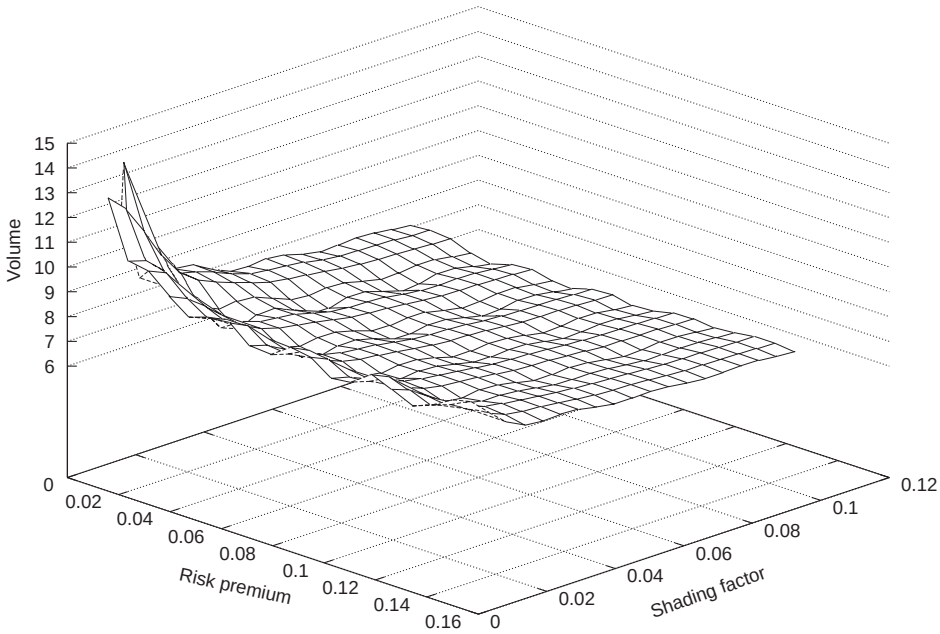

Fig. 7. Traded volume for different combinations of Δ and Π.

It is conceivable that more volatile markets can trigger further adjustments of the bid–ask spread demanded by agents, leading to a self-reinforcing mechanism that might explain volatility bursts and other similar effects. Related mechanisms are presented for example in Ref. [8], where the agents' demand is inversely proportional to the local volatility of the price, and in Ref. [12], where agents update their threshold for action on the basis of the most recent observed absolute return. We do not explore this avenue of research in the paper, but the assumption that the bid–ask spread demanded by the agents is positively correlated with the volatility of returns appears reasonable and would indeed produce a positive-feedback loop: an increment of the volatility pushes the agents to upwardly revise their spread, which in turn ignites a greater volatility. We conjecture that this potentially unstable dynamics would eventually be controlled by a reduction in the liquidity of the market, because budget constraints make impossible to sustain a trend indefinitely.

Finally, consider Fact 5. Fig. 7 shows the average traded volume per session as a function of Δ and Π. As discussed above, the traded volume decreases when the agents demand a higher bid–ask spread, because this produces a wider average gap between bids and asks which is more difficult to close.

We close by noting a fine point[2] about our comparative statics exercises. They are carried out with respect to changes in Δ and Π. Increasing Δ and Π simultaneously raise both the mean and the variance of the uniform distributions from which $\delta_i(t)$

[2]We thank one of our anonymous referees for raising it.

and π_i are extracted. Since, in particular, a higher variance implies that the agents' population tend to be more heterogeneous with respect to the prices at which they are willing to exchange the risky asset, we cannot rule out confounding effects. In order to test this, we have run separate tests using, respectively, only upward shifts (to alter the mean without affecting the variance) and mean-preserving spreads (to alter the variance without affecting the mean) of the supports of $\delta_i(t)$ and π_i. Although the results are broadly similar to those reported here, it turns out that they are less clearcut and, in some instances, monotonicity may occasionally fail.

4. Conclusions

This paper continues the study of the agent-based order-driven market originally presented in a simpler form in Ref. [1], with special emphasis on the conjoint effect of the risk premia and of the efficacy–immediacy tradeoff on the general properties of the time series of the returns. We show that even die-hard fundamentalist trading can generate fat-tailed and leptokurtic returns when it takes place in an order-driven market. This results is shown to be robust for a wide range of risk premia and intensity of efficacy in order placing. Moreover, by simulating different breeds of agents, we detect monotone relationships between the bid–ask spread demanded by the agents and some important statistical features of the returns like excess kurtosis, tail exponents and risk measures. We leave to future research the study of feedback loops affecting the spread that could explain the occurrence of volatility bursts.

Acknowledgements

We thank the participants to the "Computational Finance Workshop 2003" in Venice for useful discussions and two anonymous referees for their comments. Financial support from MIUR is acknowledged.

References

[1] M. LiCalzi, P. Pellizzari, Fundamentalists clashing over the book: a study of order-driven stock markets, Quant. Finance 3 (2003) 470–480.
[2] P. Bak, M. Paczuski, M. Shubik, Price variations in a stock market with many agents, Physica A 246 (1997) 430–453.
[3] S. Maslov, Simple model of a limit order-driven market, Physica A 278 (2000) 571–578.
[4] L. Matassini, F. Franci, On financial markets trading, Physica A 289 (2001) 526–542.
[5] G. Bottazzi, G. Dosi, I. Rebesco, Institutional architectures and behavioural ecologies in the dynamics of financial markets, J. Math. Econ. 41 (2004) 197–228.
[6] J.-P. Bouchaud, M. Mézard, M. Potters, Statistical properties of stock order books: empirical results and models, Quant. Finance 2 (2002) 251–256.
[7] C. Chiarella, G. Iori, A simulation analysis of the microstructure of double auction markets, Quant. Finance 2 (2002) 346–353.

[8] M. Raberto, S. Cincotti, S.M. Focardi, M. Marchesi, Agent-based simulation of a financial market, Physica A 299 (2001) 319–327.
[9] R. Cont, J.P. Bouchaud, Herd behaviour and aggregate fluctuations in financial markets, Macroecon. Dyn. 4 (2000) 170–196.
[10] J.D. Farmer, S. Joshi, The price dynamics of common trading strategies, J. Econ. Behav. Organ. 49 (2002) 149–171.
[11] G. Iori, A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions, J. Econ. Behav. Organ. 49 (2002) 269–285.
[12] F. Ghoulmie, R. Cont, J.-P. Nadal, Heterogeneity and feedback in an agent-based market model, J. Phys.: Condensed Matter 17 (2005) 1259–1268.
[13] D.K. Gode, S. Sunder, Allocative efficiency of markets with zero intelligence traders, J. Political Economy 101 (1993) 119–137.
[14] D. Cliff, J. Bruten, Minimal-intelligence agents for bargaining behaviors in market-based environments, Technical report, HP Laboratories Bristol, HPL-1997-91, 1997.
[15] D. Cliff, Evolution of market mechanism through a continuous space of auction-types II: two-sided auction mechanisms evolve in response to market shocks, Technical report, HP Laboratories Bristol, HPL-2002-128, 2002.
[16] R. Mehra, E.C. Prescott, The equity premium: a puzzle, J. Monetary Econ. 15 (1985) 145–161.
[17] P. Artzner, F. Delbaen, J.-M. Eber, D. Heath, Coherent measures of risk, Math. Finance 9 (1999) 203–228.