

Heterogeneity and feedback in an agent-based market model

François Ghoulmie¹, Rama Cont² and Jean-Pierre Nadal¹

¹ Laboratoire de Physique Statistique, Ecole Normale, Supérieure 24 rue Lhomond,
75231 Paris Cedex, France

² CMAP-Ecole Polytechnique, F 91128 Palaiseau, France

E-mail: ghoulmie@santafe.edu

Received 3 September 2004, in final form 17 November 2004

Published 24 March 2005

Online at stacks.iop.org/JPhysCM/17/S1259

Abstract

We propose an agent-based model of a single-asset financial market, described in terms of a small number of parameters, which generates price returns with statistical properties similar to the stylized facts observed in financial time series. Our agent-based model generically leads to the absence of autocorrelation in returns, self-sustaining excess volatility, mean-reverting volatility, volatility clustering and endogenous bursts of market activity non-attributable to external noise. The parsimonious structure of the model allows the identification of feedback and heterogeneity as the key mechanisms leading to these effects.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The study of statistical properties of financial time series has revealed a wealth of universal stylized facts which seem to be common to a wide variety of markets, instruments and periods [6, 15]. Agent-based market models, which are based on a stylized description for the behaviour of agents, attempt to explain the origins of the observed behaviour of market prices in terms of simple behavioural rules of market participants: in this approach a financial market is modelled as a system of heterogeneous, interacting agents and several examples of such models have been shown to generate price behaviour with statistical properties similar to those observed in real markets [1, 4, 5, 13, 17, 19–21, 18, 12, 26]. However most agent-based models are formulated in a complex manner and, due to their complexity, it is often not clear *which* aspect of the model is responsible for generating the stylized facts and whether all the ingredients of the model are indeed required for explaining empirical observations. This complexity also diminishes the explanatory power of such models.

We report here our findings [8] on a parsimoniously parameterized agent-based model of a single-asset financial market, which generates returns with statistical properties similar to the stylized facts observed in financial time series. Our agent-based model generically leads to the absence of autocorrelation in returns, self-sustaining excess volatility, volatility clustering and endogenous bursts of market activity non-attributable to external noise. The parsimonious structure of the model allows the identification of *heterogeneity* of strategies and *feedback* generated by the price impact of order flow as the key mechanisms leading to these effects. In particular, we show that direct interaction or herding effects are not needed to generate stylized facts. Our model presents an example where heterogeneity is endogenous and follows a stochastic evolution in time, leading to a dynamic disordered system where the disorder develops with time in a history dependent manner.

2. Statistical properties of asset returns

Time series of stock returns exhibit interesting statistical features which seem to be common to a wide range of markets and time-periods [6]:

- (i) *Excess volatility*. By ‘excess volatility’ one refers to the observation that the level of variability in market prices is much higher than can be expected based on the variability of fundamental economic variables [25] and the occurrence of large (negative or positive) returns is not always explainable by the arrival of new information on the market [25, 9].
- (ii) *Heavy tails*. The (unconditional) distribution of daily and hourly returns displays a heavy tail with positive excess kurtosis.
- (iii) *Absence of autocorrelations in returns*. (Linear) autocorrelations of asset returns are often insignificant, except for very small intraday timescales (≈ 20 min) for which microstructure effects come into play.
- (iv) *Volatility clustering*. As noted by Mandelbrot [22], ‘large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes’. In quantitative terms, while returns themselves are uncorrelated, absolute returns $|r_t(\Delta)|$ display a positive, significant and slowly decaying autocorrelation function: $\text{corr}(|r_t|, |r_{t+\Delta}|) > 0$ for Δ ranging from a few minutes to a several weeks.
- (v) *Volume/volatility correlation*. Trading volume is positively correlated with market volatility.

The fact that these empirical properties are common to a wide range of markets and time periods suggests that their origin can be retraced to some simple market mechanisms, common to many markets and thus largely independent of their ‘microstructure’ [5]. This is the basis for the development of agent-based market models, which are based on a stylized description for the behaviour of agents, and attempt to explain the origins of the observed behaviour of market prices as emerging from simple behavioural rules of a large number of heterogeneous market participants. Several examples of such models have been shown to generate price behaviour with statistical properties similar to those observed in real markets [1, 3, 4, 17, 19, 21, 18, 12, 26]. Numerical simulations of many of the models above lead to time series of ‘returns’ which have properties consistent with (some of) the empirical stylized facts observed above. However, due to the complexity of such models it is often not clear *which* aspect of these models is responsible for generating the stylized facts and whether all the ingredients of the model are indeed required for explaining empirical observations. This complexity also diminishes the explanatory power of such models.

A key point in [7] which leads to heavy tails in the distribution of order flow is to allow for *investor inertia*—the fact that most market participants trade very infrequently. A possible

mechanism for generating investor inertia is *threshold response* in the behaviour of market participants [14]: the risk aversion of agents which leads them to be inactive if uncertain about their action. Based on these remarks, we formulate a model of a single-asset market retaining the ingredients above.

3. Description of the model

Our model describes a market where a single asset, whose price is denoted by p_t , is traded by N agents. Trading takes place at discrete periods $t = 0, 1, 2, \dots$. We will see that, provided the parameters of the model are chosen in a certain range, we will be able to interpret these periods as ‘trading days’. At each period, agents have the possibility of sending an order to the market for buying or selling a unit of asset: denoting by $\phi_i(t)$ the demand of the agent, we have $\phi_i(t) = 1$ for a buy order and $\phi_i(t) = -1$. We allow the value $\phi_i(t)$ to be zero; the agent is then inactive at period t . The inflow of public information is modelled by a sequence of IID Gaussian random variables $(\epsilon_t, t = 0, 1, 2, \dots)$ with $\epsilon_t \sim N(0, D^2)$. ϵ_t represents the value of a common signal received by all agents at date $t - 1$. The signal ϵ_t is a forecast of the future return r_t and each agent has to decide whether the information conveyed by ϵ_t is significant, in which case she will place a buy or sell an order according to the sign of ϵ_t .

The trading rule of each agent $i = 1, \dots, N$ is represented by a (time-varying) decision threshold $\theta_i(t)$. The threshold $\theta_i(t)$ can be viewed as the agent’s (subjective) view on volatility. The trading rule we study may be seen as a stylized example of threshold behaviour: without sufficient external stimulus ($|\epsilon_t| \leq \theta_i(t)$), an agent remains inactive $\phi_i(t) = 0$ and if the external signal is above a certain threshold, the agent will act: if $\epsilon_t > \theta_i(t)$, $\phi_i(t) = 1$, if $\epsilon_t < -\theta_i(t)$, $\phi_i(t) = -1$. The corresponding demand generated by the agent is therefore given by:

$$\phi_i(t) = 1_{\epsilon_t > \theta_i(t)} - 1_{\epsilon_t < -\theta_i(t)}. \quad (1)$$

The excess demand is then given by $Z_t = \sum_{i=1}^N \phi_i(t)$. A non-zero value of Z produces a change in the price given by

$$r_t = \ln \frac{p_t}{p_{t-1}} = g\left(\frac{Z_t}{N}\right) \quad (2)$$

where the price impact function $g : \mathbb{R} \mapsto \mathbb{R}$ is an increasing function with $g(0) = 0$. We define the (normalized) market depth λ by: $g'(0) = 1/\lambda$. While most of the analysis below holds for a general price impact function g , in some cases it will be useful to consider a linear price impact: $g(z) = z/\lambda$.

Initially, we start from a population distribution F_0 of thresholds: $\theta_i(0), i = 1 \dots N$ are positive IID variables drawn from F_0 . Updating of strategies is *asynchronous*: at each time step, any agent i has a probability $0 \leq s \leq 1$ of updating her threshold $\theta_i(t)$. Thus, in a large population, s represents the fraction of agents updating their views at any period; $1/s$ represents the typical time period during which an agent will hold a given view $\theta_i(t)$. If periods are to be interpreted as days, s is typically a small number $s \simeq 10^{-1} - 10^{-3}$. When an agent updates her threshold, she sets it to be equal to the recently observed absolute return, which is an indicator of recent volatility $|r_t| = \left| \ln \frac{p_t}{p_{t-1}} \right|$. Introducing IID random variables $u_i(t), i = 1 \dots N, t \geq 0$ uniformly distributed on $[0, 1]$, which indicate whether agent i updates her threshold or not:

$$\theta_i(t) = 1_{u_i(t) < s} |r_t| + 1_{u_i(t) \geq s} \theta_i(t - 1). \quad (3)$$

This way of updating can be seen as a stylized version of various estimators of volatility based on moving averages of absolute or squared returns. It is also corroborated by a recent empirical

study by Zovko and Farmer [27], who show that traders use recent volatility as a signal when placing orders.

The asynchronous updating scheme proposed here avoids introducing an artificial ordering of agents as in sequential choice models. The random nature of updating is also a parsimonious way to introduce heterogeneity in timescales, a feature believed to be important [19], without introducing extra parameters. Given this random updating scheme, even if we start from an initially homogeneous population $\theta_i(0) = \theta_0$, heterogeneity creeps into the population through the updating process and evolves in a random manner, leading to a history-dependent disordered system.

Let us recall the main ingredients of the model. At each time period:

- (i) Agents receive a common signal $\epsilon(t) \sim N(0, D^2)$.
- (ii) Each agent i compares the signal to her threshold $\theta_i(t)$.
- (iii) If $|\epsilon(t)| > \theta_i(t)$ the agent considers the signal as significant and generates an order $\phi_i(t)$ according to (1).
- (iv) The market price is impacted by the excess demand and moves according to (2).
- (v) Each agent updates, with probability s , her threshold according to (3).

Compared to most agent-based models considered in the literature, there is no exogenous ‘fundamental price’ process and we do not distinguish between ‘fundamentalist’ and ‘chartist’ traders. Also, the same information is available to all agents but they differ in the way they *process* the information. We do not introduce any ‘social interaction’ among agents: no notion of locality, lattice or graph structure is introduced. The model has very few parameters: s describes the average updating frequency, D the standard deviation of the noise representing the news arrival process, the market depth λ and the number of agents N which is typically large. We will observe nevertheless that this simple model generates time series of returns with interesting properties similar to empirically observed properties of asset returns.

4. Simulation results

In order for a direct comparison with empirical stylized facts to be meaningful, we have to consider that in the case of empirical data only a single sample path of the price is available and (unconditional) moments are computed by averaging over the (single) sample path. We therefore adopt a similar approach here: after simulating a sample path of the price p_t for $T = 10^4$ periods, we compute the time series of returns $r_t = \ln(p_t/p_{t-1})$, $t = 1 \dots T$, their histogram, a moving average estimator of the standard deviation of returns (‘volatility’), the sample autocorrelation function of returns and the sample autocorrelation function of absolute returns. In order to decrease the sensitivity of results to initial conditions, we allow for a transitory regime and discard the first 10^3 periods before averaging.

In order to interpret the trading periods as ‘days’ and compare the results with properties of daily returns, we note that when g is linear $|r_t| \leq 1/\lambda$ and choose $5 \leq \lambda \leq 20$ which allows a (maximal) range of daily returns between 5% and 20%. Also, the amplitude D of the input noise can be chosen such as to reproduce a realistic range of values for the (annualized) volatility: this leads to choosing D in the range 10^{-3} – 10^{-2} . Let us emphasize that we are discussing the calibration of the *order of magnitude* of parameters, not fine-tuning them to a set of critical values. The results discussed in the sequel are generic within this range of parameters. Figures 1 and 2 illustrate typical sample paths obtained with different parameter values: they all generate series of returns with realistic ranges and realistic values of annualized volatility. For each series, we represent the histogram of returns both in linear and logarithmic




Figure 1. Numerical simulation of the model with updating frequency $s = 0.01$ (average updating period: 100 'days') $N = 1000$ agents, $D = 0.001$ and $\lambda = 10$.

scales, the ACF of returns C_r , the ACF of absolute returns $C_{|r|}$. The return series obtained possess regularities which match the properties outlined in section 2:

- (i) Excess volatility. The sample standard deviation of returns can be much larger than the standard deviation of the input noise representing news arrivals $\hat{\sigma}(t) \gg D$.
- (ii) Mean-reverting volatility. The market price fluctuates endlessly and the volatility, as measured by the moving average estimator $\hat{\sigma}(t)$, goes neither to zero nor to infinity and displays a mean-reverting behaviour.
- (iii) The simulated process generates a leptokurtic distribution of returns with (semi-)heavy tails, with an excess kurtosis around $\kappa \simeq 7$. As shown in the logarithmic histogram plots in figures 1, 2, the tails exhibit an approximately exponential decay, as observed in various studies of daily returns [10].
- (iv) The returns are uncorrelated. The sample autocorrelation function of the returns exhibits an insignificant value (very similar to that of asset returns) at all lags, indicate absence of linear serial dependence in the returns.




Figure 2. Numerical simulation of the model with updating frequency $s = 0.1$ (average updating period: 10 'days') $N = 1500$ agents, $D = 0.001$ and $\lambda = 10$.

- (v) Volatility clustering. The autocorrelation function of absolute returns remains significantly positive over many time lags, corresponding to persistence of the amplitude of returns a timescale $\simeq 1/s$.

5. Some limiting cases

- (i) *Feedback without heterogeneity.* In the case where $s = 1$, all agents synchronously update their threshold at each period. Consequently, the agents have the same thresholds, given by the last periods absolute return: $\theta_i(t) = |r_{t-1}|$ and will therefore generate the same order: $Z_t = N\phi_1(t) \in \{0, N, -N\}$. So, the return r_t depends on the past only through the absolute return $|r_{t-1}|$:

$$r_t = f(|r_{t-1}, \epsilon_t|) = g(N)1_{\epsilon_t > |r_{t-1}|} + g(-N)1_{\epsilon_t < -|r_{t-1}|},$$

a dependence structure typical of ARCH models [11], leading to uncorrelated returns and volatility clustering. In this case, the distribution of r_t conditional on $|r_{t-1}|$ is actually a trinomial distribution: $r_t \in \{0, g(N), g(-N)\}$, which is not realistic. Simulation studies show that a similar behaviour persists for $1 - s \ll 1$, leading to tri-modal distributions. This confirms our intuition that the updating probability s should be chosen small.

- (ii) *Heterogeneity without feedback.* In the case where $s = 0$, no updating takes places: the trading strategies, given by the thresholds θ_i , are unaffected by the price behaviour and

the *feedback* effect is no longer present. Heterogeneity is still present: the distribution of the thresholds remains identical to what it was at $t = 0$. The return r_t depends only on ϵ_t :

$$r_t = g\left(\frac{1}{N} \sum_{i=1}^N 1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}\right) = F(\epsilon_t).$$

We conclude therefore that the returns are IID random variables, obtained by transforming the Gaussian IID sequence (ϵ_t) by the nonlinear function F given in (ii), whose properties depend on the (initial) distribution of thresholds $(\theta_i, i = 1 \dots N)$. The log-price then follows a (non-Gaussian) random walk and the model does not exhibit volatility clustering.

6. Behaviour of prices and volatility

The two limiting cases above show that, in order to obtain the interesting statistical properties observed in the simulated examples shown above, it is necessary to have $0 < s \ll 1$; both feedback and heterogeneity are essential ingredients. In the general case we have the following properties:

- *Markovian dynamics.* The thresholds $[\theta_i(t), i = 1 \dots N]$ follow a Markov chain in $\{g(k), k = 0 \dots N\}$. We have $\theta_i(t + 1) = \theta_i(t)$ with probability $1 - s$ and

$$\theta_i(t + 1) = |r_t| = \left| g\left(\frac{1}{N} \sum_{i=1}^N [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}]\right) \right| \quad \text{with probability } s.$$

In fact given that agents are indistinguishable and only the empirical distribution of threshold values affects the returns, defining $N_k(t) = \sum_{i=1}^N 1_{[0, a_k] }(\theta_i(t))$ then $(N_k(t), k = 0 \dots N - 1)_{t=0,1,\dots}$ evolves as a Markov chain in $\{0, \dots, N\}^N$. $N(t) = (N_k(t), k = 0 \dots N - 1)$ is none other than the (cumulative) population distribution of the thresholds. The fact that $N(t)$ itself follows a Markov chain means that the population distribution of thresholds is a *random measure* on $\{0, \dots, N\}$, which is characteristic of disordered systems [23], even if we start from a deterministic set of values for the initial thresholds (even identical ones). Here the disorder is endogenous and is generated by the random updating mechanism.

- *Excess volatility.* In this model, the volatility of the news arrival process is quantified by D which is the standard deviation of the external noise ϵ_t , whereas the volatility of the returns can be measured a posteriori as the (conditional or unconditional) standard deviation of r_t . As seen from the nonlinear relation between ϵ_t and r_t ,

$$r_t = g\left(\frac{\sum_{i=1}^N 1_{\epsilon_t > \theta_i(t)} - 1_{\epsilon_t < -\theta_i(t)}}{\lambda N}\right)$$

even after conditioning on the current states of agents $\theta_i(t), i = 1 \dots N$, equation (6) yields a nonlinear relation between the input noise ϵ_t and the returns which can have the effect of amplifying the noise by an order of magnitude or more. In the simulation example shown in figure 1, $D = 10^{-3}$ which corresponds to an annualized volatility of 1.6%, while the annualized volatility of returns is in the range of 20%, an order of magnitude larger; the order of magnitude of the volatility of returns may be quite different from that of the input noise.

- *Absence of autocorrelation.* From the dynamic equations of the model

$$Z_t = \frac{1}{N} \sum_{i=1}^N \phi_i(t) = \frac{1}{N} \sum_{i=1}^N [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}] \tag{4}$$

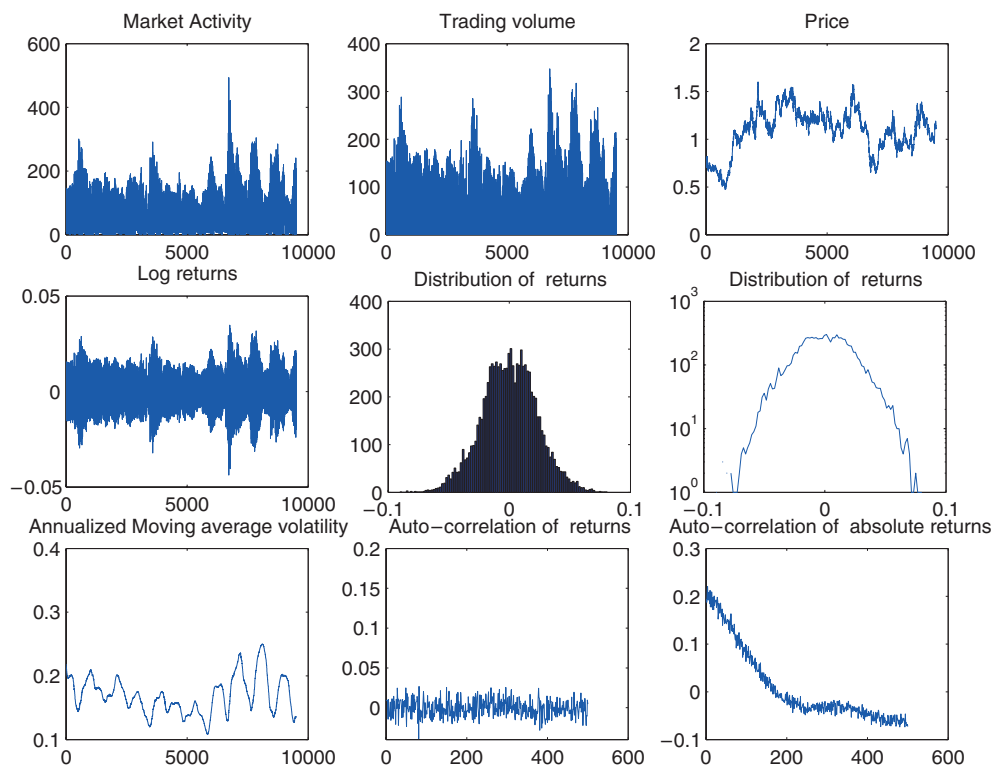


Figure 3. Left: correlation timescale τ_c of absolute returns, as a function of the updating period $1/s$. Right: evolution of the portfolio of a typical agent, with long periods of inactivity punctuated by bursts of activity.

$$r_t = g(Z_t) = g\left(\frac{1}{N} \sum_{i=1}^N [1_{\epsilon_i > \theta_i} - 1_{\epsilon_i < -\theta_i}]\right) \quad (5)$$

one can deduce that, if g is an odd function (in particular if g is linear), then asset returns $(r_t)_{t \geq 0}$ are uncorrelated; $\text{cov}(r_t, r_{t+1}) = 0$. This is due to the fact that the trading/nontrading decision is based only on the amplitude of the signal, not its sign. The sign of the return is determined by the sign of the common signal, which is independent across periods.

- *Investor inertia.* Except in times of crisis or market crash, at a given point in time only a small proportion of stockholders are actually trading in the market. As a result, the (daily) order flow for a typical stock can be much smaller than the market capitalization. This phenomenon, sometimes referred to as *investor inertia*, is a generic outcome in our model due to threshold behaviour of agents. Starting from an initial holding of $\pi_i(0)$, the quantity of asset held by agent i is given by $\pi_i(t) = \sum_{\tau=0}^t \phi_i(\tau)$. Figure 3 displays the evolution of the portfolio $\pi_i(t)$ of a typical agent; short periods of activity (trading) are separated by long periods of inertia, where the portfolio remains constant. This ‘inertia’ increases in periods of high volatility, an effect similar to the behaviour of risk-averse agent.
- *Clustering and mean-reversion in volatility.* Many market microstructure models—especially those with learning or evolution—converge over large time intervals to an

equilibrium where prices and other aggregate quantities cease to fluctuate randomly. By contrast, in the present model, prices fluctuate endlessly and the volatility exhibits mean-reverting behaviour. Suppose we are in a period of 'low volatility'; the amplitude $|r_t|$ of returns is small. Agents who update their thresholds will therefore update them to small values, become more sensitive to news arrivals, thus generating higher excess demand and thus increasing the amplitude of returns. Conversely, in a period of high volatility, agents will update their threshold values to high values and become less reactive to the incoming signal: this increase in investor inertia will thus decrease the amplitude of returns. The mean reversion time in the volatility is therefore the time it takes for agents to adjust their thresholds to current market conditions, which is of the order of $\tau_c = 1/s$.

When the amplitude of the noise is small it can be shown [8] that volatility decays exponentially in time and increases through upward 'jumps'. This behaviour is actually similar to that of a class of stochastic volatility models introduced by Barndorff-Nielsen and Shephard [2] and successfully used to describe various econometric properties of returns.

7. Conclusion

We have presented a parsimonious agent-based model capable of reproducing the main empirical stylized facts described in section 2, based on three main ingredients:

- (i) Threshold behaviour of agents.
- (ii) Heterogeneity of agent strategies, generated endogenously through random asynchronous updating of thresholds.
- (iii) Feedback of recent price behaviour on agents behaviour.

Numerical simulations of the model generically produce time series that capture the stylized facts observed in asset returns. Due to the simple structure of the model, these simulation results can be explained by a theoretical analysis of the price process in the model. These observations illustrate that these three ingredients suffice for reproducing several empirical stylized facts such as heavy tails, absence of autocorrelation in returns and volatility clustering, with realistic values in the timescales involved and without any exogenous 'fundamental' price, direct interaction between agents or distinction between 'chartist' or 'fundamentalist' traders. These results question some previous conclusions regarding the origins of stylized properties of asset returns previously drawn from simulation of agent-based models and call for a closer, critical look at this issue through the study of a wider variety of agent-based market designs. These points are further developed in [8].

References

- [1] Arthur W, Holland J, LeBaron B, Palmer J and Tayler P 1997 Asset pricing under endogeneous expectations in an artificial stock market *The Economy as an Evolving Complex System* vol 2, ed W Arthur, S Durlauf and D Lane (Reading, MA: Perseus Books) pp 15–44
- [2] Barndorff-Nielsen O E and Shephard N 2001 Non-Gaussian Ornstein–Uhlenbeck based models and some of their uses in financial econometrics *J. R. Stat. Soc. B* **63** 167–241
- [3] Brock W and Hommes C 1998 Heterogeneous beliefs and route to chaos in a simple asset pricing model *J. Econ. Dyn. Control* **22** 468–85
- [4] Challet D, Chessa A, Marsili M and Zhang Y 2001 From minority games to real markets *Quant. Finance* **1** 168–76
- [5] Cont R 1998 Modeling economic randomness: statistical mechanics of market phenomena *Statistical Physics in the 21st Century: The James B McGuire Festschrift* ed M Batchelor and L Wille (Singapore: World Scientific)

-
- [6] Cont R 2001 Empirical properties of asset returns: stylized facts and statistical issues *Quant. Finance* **1** 1–14
 - [7] Cont R and Bouchaud J-P 2000 Herd behaviour and aggregate fluctuations in financial markets *Macroecon. Dyn.* **4** 170–95
 - [8] Cont R, Ghoulmie F and Nadal J-P 2004 Threshold behaviour and volatility clustering in financial markets *Working Paper*
 - [9] Cutler D, Poterba J and Summers L 1989 What moves stock prices? *J. Portfolio Manage.* 4–12
 - [10] Ding Z, Granger C and Engle R 1983 A long memory property of stock market returns and a new model *J. Empirical Finance* **1** 83–106
 - [11] Engle R 1995 *ARCH Models* (Oxford: Oxford University Press)
 - [12] Farmer J D and Joshi S 2002 The price dynamics of common trading strategies *J. Econ. Behav. Organ.* **49** 149–71
 - [13] Giardina I and Bouchaud J-P 2003 Bubbles, crashes and intermittency in agent based market models *Eur. Phys. J. B* **31** 421–37
 - [14] Granovetter M 1978 Threshold models of collective behaviour *Am. J. Sociol.* **83** 1420–43
 - [15] Guillaume D, Dacorogna M, Davé R, Müller U, Olsen R and Pictet O 1997 From the birds eye view to the microscope: a survey of new stylized facts of the intraday foreign exchange markets *Finance and Stochastics* **1** 95–131
 - [16] Hommes C H, Gaunersdorfer A and Wagener F O 2003 Bifurcation routes to volatility clustering under evolutionary learning *Technical Report*, CenDEF
 - [17] Iori G 2002 A microsimulation of traders activity in the stock market: the role of heterogeneity, agents interactions and trade frictions *J. Econ. Behav. Organ.* **49** 269–85
 - [18] Kirman A and Teyssiere G 2002 Microeconomic models for long-memory in the volatility of financial time series *Stud. Nonlinear Dyn. Econometrics* **5** 281–302
 - [19] LeBaron B 2001 Evolution and time horizons in an agent-based stock market *Macroecon. Dyn.* **5** 225–54
 - [20] LiCalzi M and Pellizzari P 2003 Fundamentalists clashing over the book: a study of order-driven stock markets *Quant. Finance* **3** 470–80
 - [21] Lux T and Marchesi M 2000 Volatility clustering in financial markets *Int. J. Theor. Appl. Finance* **3** 675–702
 - [22] Mandelbrot B B 1963 The variation of certain speculative prices *J. Business* **36** 392–417
 - [23] Mézard M, Parisi G and Virasoro M 1989 *Spin Glasses and Beyond* (Singapore: World Scientific)
 - [24] Orlean A 1995 Bayesian interactions and collective dynamics of opinions *J. Econ. Behav. Organ.* **28** 257–74
 - [25] Shiller R 2000 *Irrational Exuberance* (Princeton, NJ: Princeton University Press)
 - [26] Stauffer D 2001 Percolation models of financial market dynamics *Adv. Complex Syst.* **4** 19–27
 - [27] Zovko I and Farmer J D 2002 The power of patience: a behavioural regularity in limit order placement *Quant. Finance* **2** 387–92