Heterogeneity and feedback in an agent-based market model

François GHOULMIE[a,b,1], Rama CONT[a], Jean-Pierre NADAL[b]

(a)Centre de Mathématiques Appliquées
Ecole Polytechnique, F-91128 Palaiseau, France.
(b)Laboratoire de Physique Statistique
Ecole Normale Supérieure, F-75321 Paris Cedex 05, France

July 1, 2004

DRAFT

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. We show that the joint effect of feedback and heterogeneity leads to a market price which fluctuates endlessly and a volatility which displays a mean-reverting behavior : the volatility goes neither to zero nor to infinity in the long-run.
Our agent-based model generically leads to absence of autocorrelation in returns, stochastic volatility, excess volatility, volatility clustering and endogeneous bursts of market activity non-attributable to external noise. The parsimonious structure of the model allows to identify the mechanism leading to these effects. We investigate theoretically some properties of this model and present extensive numerical simulations of other properties.

Keywords: agent-based model, financial markets, stylized facts, heterogeneity, feedback, asynchronous updating.

1

[1]Corresponding author. Tel.: +33 1 69 33 47 48; fax: +33 1 69 33 30 11.
E-mail adress: ghoulmie@cmapx.polytechnique.fr (F. Ghoulmié).
ghoulmie@santafe.edu.
This paper was prepared for the 10th International Conference of the Society for Computational Economics, July 2004, Amsteram, The Netherlands.

Contents

The study of statistical properties of financial time series has revealed a wealth of interesting stylized facts which seem to be common to a wide variety of markets, instruments and periods. Market microstructure models, which delve into the details of agent preferences and their interactions, may be useful for normative design of market mechanisms but fail to explain why such stylized facts are common to markets with such different microstructures.
Agent-based market models, which are based on a stylized description for the behavior of agents, attempt to explain the origins of the observed behavior of market prices in terms of simple behavioral rules of market participants: in this approach a financial market is modeled as a system of heterogeneous, interacting agents and several examples of such models have been shown to generate price behavior similar to those observed in real markets.
Agent-based models studied in the literature have pointed to various possible explanations for empirical stylized facts: herd behavior, social interaction and mimetism, heterogeneity, investor inertia and switching between "chartist" and "fundamentalist" behavior have been invoked as possible mechanisms. However most of these 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 propose here a parsimoniously parametrized 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 absence of autocorrelation in returns, mean-reverting stochastic volatility, excess volatility, volatility clustering and endogeneous bursts of market activity non-attributable to external noise. The parsimonious structure of the model allows to identify *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. While heterogeneity of preferences and endowments has also been studied in more classical frameworks in microeconomics, our model presents an example where heterogeneity is endogeneous and follows a stochastic evolution in time.
The article is structured as follows. Section 1 recalls some stylized empirical facts about returns of financial assets and reviews some agent-based models presented in the literature to explain the stylized facts. Section 2 presents our model; simulation results are presented in section 3 and a theoretical discussion is given in section 4 and section 5.

1 Stylized properties of asset returns: facts and models.

1.1 Stylized statistical propeties of asset returns.

Time series of asset returns exhibit interesting statistical features which seem to be common to a wide range of markets and time-periods as reviewed by Cont (2001):

• **Excess volatility:** many empirical studies point out to the fact that it is difficult to justify the observed level of variability in asset returns by variations in "fundamental" economic variables. In particular, the occurence of large (negative or positive) returns is not always explainable by the arrival of new information on the market as noted by Cutler et al. (1989), and Shiller (2000).

• **Heavy tails:** the (unconditional) distribution of returns displays a heavy tail with positive excess kurtosis.

• **Absence of autocorrelation in returns:** (linear) autocorrelations of asset returns are often insignificant, except for very small intraday time scales ($\simeq$20 minutes) for which microstructure effects come into play.

• **Volatility clustering:** as noted by Mandelbrot (1963), "large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes". A quantitative manifestation of this fact is that, while returns themselves are uncorrelated, absolute returns $|r_t(\Delta)|$ or their squares display a positive, significant and slowly decaying autocorrelation function: $corr(|r_t|, |r_{t+\Delta|}) > 0$ for Δ ranging from a few minutes to several weeks.

• **Volume/volatility correlation:** trading volume is positively correlated with market volatility.

These stylized facts are model-free, "nonparametric" statements on properties of returns. While they do not pinpoint a single stochastic process as a candidate for the price dynamics, the fact that they are common to a wide variety of markets and periods has intrigued many researchers who have sought to explain their origin by relating them to the behavior of market agents.

1.2 Agent-based models of financial markets.

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". This is the basis for the development of agent-based market models, which are based on a stylized description for the behavior of agents and attempt to explain the origins of the observed behavior of market prices as emerging from simple behavioral rules of a large number of heterogeneous market participants. Stauffer (2001) proposed for example percolation models that generate price behavior with statistical properties similar to those observed in real markets.

Agent-based models studied in the literature have pointed to various possible explanation for empirical stylized facts. Mimetism (e.g., Orléan, 1995) and herd behavior (e.g., Cont & Bouchaud, 2000 and Stauffer et al., 1999) has been suggested as a possible explanation for heavy tails in returns and the occurence of bubbles (e.g., Banerjee, 1993).

Another mechanism which has been considered is endogeneous switching between "chartist" and "fundamentalist" behavior as studied by Lux & Marchesi (2000) and Kirman & Teyssiere (2002). In such models one starts from an exogeneous "fundamental value" for an asset and investors can switch between a strategy based on the fundamental value or a trend-following rule. This switching is generated by the fact that backward-looking investors compare the benefits of having used one strategy or the other in the recent past, leading to boom-bust cycles in prices. Brock & Hommes (1998) study deterministic variants and show that such behavior can also lead to chaotic behavior which mimicks some "statistical" stylized facts. Similarly, Gaunersdorfer (2000) studies a model in which traders have heterogeneous expectations concerning future prices and update their beliefs according to a risk adjusted performance measure and to market conditions: in the case where agents can only choose between two different predictors chaotic behavior ensues. Chiarella et al. (2002) consider a wider spectrum of forecasting strategies to study the effects of heterogeneous beliefs and attitudes towards risk on the dynamics of asset prices and wealth. Some models enrich this approach by adding the role of a market maker.

Evolutionary variants of these models where "natural" selection favors the survival of best performing strategies have been studied by Arthur et al. (1997), Blume & Easley (1990), Farmer & Joshi (2002), Hommes et al. (2003).

Heterogeneity in agents' time scale are also believed to be responsible for a number of stylized facts. Long term traders naturally focus on long-term behavior of prices neglecting fluctuations at the smallest time scale, whereas short-term traders are not concerned with price movements on the long-run but rather aim to exploit short-term predictability. The effects of the diversity in time horizons on price dynamics have been studied by LeBaron (2001) in an artificial stock market. He concluded that the presence of heterogeneity in horizons may lead to an increase in return variability, as well as volatility-volume relationships similar to those of actual markets.

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 Cont & Bouchaud (2000) work which leads to heavy tails in the distribution of order flow is to allow for *investor inertia*-the fact that most market partcipants trade very infrequently. Horst et al. (2003) have also shown that this ingredient leads to long range dependence in returns; note however that such "long range dependence", if any, is actually observed in volatility whereas returns are believed to be uncorrelated. Investor inertia will also be a key feature of our model; however, instead of postulating investor inertia as a starting point, we will define a mechanism which leads to it as a consequence.
A possible mechanism for generating investor inertia is threshold behavior. Threshold response in the behavior of market participants can be seen either as resulting from trade friction or, more fundamentally, from the risk aversion of agents which leads them to be inactive if uncertain about their action. There are also trading rules practiced by chartists such as John Bollinger rule using threshold response. In a pioneering work, Granovetter (1983) has suggested threshold behavior in the behavior of individuals as the possible origin for collective phenomena and large fluctuations in aggregate quantities. We will use threshold behavior as a crucial ingredient in formulating our market model. Based on these remarks, we now formulate a model of a single-asset market retaining the ingredients above.

2 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 dates $t = 0, 1, 2, ..$[2]. At each period, every agent receives public news about the asset's performance and, using a subjective criterion, judge whether this news is significant. If the news is found to be significant, the agent places a buy or sell order, depending on whether the news received is pessimistic or optimistic. Prices then moved up or down according to excess demand. We now describe these ingredients in more precise terms.

2.1 Trading rules.

At each period, agents have the possibility to send 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 $\phi_i(t) = -1$ for a sell order. We allow the value $\phi_i(t)$ to be zero; the agent is then inactive at period t. The inflow of new information is modeled by a sequence of IID Gaussian random variables $(\epsilon_t, t = 0, 1, 2, ..)$ with $\epsilon_t \sim N(0, D^2)$. ϵ_t represents the value of a common signal received by all agents at date t. The signal ϵ_t is a forecast of the future return

[2] 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".

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 order according to the sign of ϵ_t.
The trading rule of each agent $i = 1, ..., N$ is represented by a (time-varying) decision threshold $\theta_i(t)$. The threshold $\theta_i(t)$ can be viewed as the agents (subjective) view on volatility. By comparing the signal to her threshold, the agent decides whether the news is significant enough to generate a trade ($|\epsilon_t| > \theta_i(t)$):

$$\begin{aligned} \text{if} \quad & \epsilon_t > \theta_i^+, \ \phi_i = 1 \\ \text{if} \quad & \epsilon_t < \theta_i^-, \ \phi_i = -1 \\ \text{otherwise} \quad & \phi_i = 0. \end{aligned} \tag{1}$$

This trading rule may be seen as a stylized example of threshold behavior: without sufficient external stimulus, an agent remains inactive and if the external signal is above a certain threshold, the agent will act. The corresponding demand generated by the agent is therefore given by:

$$\phi_i(t) = 1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}. \tag{2}$$

2.2 Price response to aggregate demand.

The aggregate excess demand is then given by:

$$Z_t = \sum_i \phi_i(t). \tag{3}$$

A non zero value of Z_t produces a change in the price, and the resulting log return is given by :

$$r_t = \ln \frac{p_t}{p_{t-1}} = g(\frac{Z_t}{N}). \tag{4}$$

where the price impact function g : $\Re \to \Re$ is an increasing function with g(0)=0. We define the (normalized) market depth λ by:

$$g'(0) = 1/\lambda. \tag{5}$$

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: $g(z) = z/\lambda$.

2.3 Updating of strategies.

As we noted above, the threshold $\theta_i(t)$ represents the view of agent i on recent market volatility: these thresholds are updated by agents from time to time to reflect the amplitude of recent returns. Initially, we start from a population distribution F_0 of thresholds: $\theta_i(0), i = 1..N$ are positive IID variables drawn from F_0.

Updating of strategies is *asynchronous*: at each time step, any agent i has a probability s ($s \in [0,1]$) 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 $|r_t| = |\ln \frac{p_t}{p_{t-1}}|$. Introducing IID random variables $u_i(t), i = 1..N, t \geq 0$ uniformly distributed on $[0,1]$, which indicate whether agent i updates her threshold or not, we can write the updating rule as:

$$\theta_i(t) = 1_{u_i(t)<s}|r_t| + 1_{u_i(t)\geq s}\theta_i(t-1). \tag{6}$$

Here ϵ_t represents randomness due to public news arrivals whereas the random variables $u_i(t)$ represent idiosyncratic sources of randomness. This way of updating can be seen as a stylized version of various estimators of volatility based on moving averages or squared returns. It is also corroborated by a recent empirical study by Zovko & Farmer (2002), 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 (e.g, Banerjee, 1993). The random nature of updating is also a parsimonious way to differentiate agents through their updating "frequencies", feature believed to be important (e.g., LeBaron, 2001). Hong et al. (2000) present empirical evidence supporting the view that inexperienced analysts revise their forecasts more frequently than inexperienced analysts; more generally, existence of agents with various time scales of intervention has been stressed in many studies of agent-based models and this random updating allows to introduce heterogeneity in time scales without introducing extra parameters.
Note that, 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. In this sense, the heterogeneity of agents strategies is endogeneous in this model and, as we will see below, evolves in a random manner.

2.4 Summary.

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

•agents receive a common signal $\epsilon_t \sim N(0, D^2)$.

•each agent i compares the signal to her threshold $\theta_i(t)$.

•if $|\epsilon_t| > \theta_i(t)$ the agent considers the signal as significant and generates an order $\phi_i(t)$ according to (1).

•The market price is impacted by the excess demand and moves according to (4).

•each agent updates, with probability s, her threshold according to (6).

With regard to some of the agent-based models considered in the literature, some important aspects are the following:

•There is no exogeneous "fundamental price" process: prices move through market fluctuations of supply and demand. In particular, we do not distinguish between "fundamentalist" and "chartist" traders.

•No information asymetry: the same information is available to all agents. Agents differ in the way they *process* the information.

•Absence of "social interaction": agents interacts indirectly via the price, as in standard Walrasian markets. We do not introduce any "social interaction" among agents. In particular, no notion of locality, lattice or graph structure is introduced.

•Endogeneous heterogeneity: heterogeneity of agents behavioral rules appears endogeneously due to the asynchronous updating sheme.

The model has very few parameters: s describes the average updating frequency, D the standard deviation of the noise representing the news arrival process and λ the market depth. Furthermore, as we will observe in the next section, if we require to interpret the trading periods as "days" this will put a further restriction on the parameters, reducing the effective number of parameters. We will observe nevertheless that this simple model generates time series of returns with interesting dynamics and properties similar to empirically observed properties of asset returns.

3 Numerical simulations

The model described above is straightforward to simulate. We describe in this section the simulation procedure, describe the quantities of interest and presents some typical results. The simulation procedure allows us to identify generic properties of the model and calibrate the range of parameters in accordance with empirical data on asset returns.

3.1 Simulation procedure.

The state of the system at each time period is described by the vector $(\theta_i(t), i = 1..N)$ of thresholds. The thresholds $\theta_i(0)$ are initialized by drawing from an IID distribution F_0. Simulation is done through an iterative procedure, each iteration repeating the steps described in Section 2.4. Although the model setting accomodates for more general price impact functions, in absence of an empirically motivated parametric form, we have chosen a linear function $g(x) = x/\lambda$. This choice can be viewed as a linearization of a more general g, valid for small values of excess demand or in a market with large market depth.
In a usual Monte Carlo simulation approach, expectations, moments and distributions of quantities of interest can be computed by running independent simulations and averaging the quantities of interest over the simulation runs. Note however that, 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 following quantities:

•the time series of returns $r_t = ln(p_t/p_{t-1}), t = 1..T$.

•the histogram of returns, which is an estimator of its unconditional distribution.

•a moving average estimator of the standard deviation of returns:

$$\hat{\sigma}^2(t) = 250[\frac{1}{T'}\sum_{t'=t-T'+1}^{t} |r_t|^2 - (\frac{1}{T'}\sum_{t'=t-T'+1}^{T'} r_t)^2]. \tag{7}$$

This quantity is a frequently used indicator for "volatility". We "annualize" it by multiplying the "daily" estimate by 250.

•the sample autocorrelation function of returns:

$$C_r(\tau) = \frac{T}{\sum_{t=1}^{T}|r_t|^2}[\frac{1}{T-\tau}\sum_{t=1}^{T-\tau} r_t r_{t+\tau} - (\frac{1}{T-\tau}\sum_{t=1}^{T-\tau} r_t)(\frac{1}{T-\tau}\sum_{t=\tau+1}^{T} r_t)] \tag{8}$$

•the sample autocorrelation function of absolute returns:

$$C_|r|(\tau) = \frac{T}{\sum_{t=1}^{T}|r_t|^2}[\frac{1}{T-\tau}\sum_{t=1}^{T-\tau} |r_t||r_{t+\tau}| - (\frac{1}{T-\tau}\sum_{t=1}^{T-\tau} |r_t|)(\frac{1}{T-\tau}\sum_{t=\tau+1}^{T} |r_t|)] \tag{9}$$

These quantities can then be used to compare with the empirical stylized facts described in Section 1.1. Finally, in order to decrease the sensitivity of results to initial conditions, we allow for an initial transitory regime and discard the first 10^3 periods before averaging. The rationale behind the choice of this initial period will be discussed below.

3.2 Choosing the range of parameters

Simulation of the model requires the specification of the parameters s, D, λ, the number of agents N and the specification of the initial distribution of the thresholds.

In order to interpret the trading periods as "days" and compare the results obtained to properties of daily returns, some restrictions must be imposed on parameter values. First, note that the one-period returns are bounded by $max\{|g(x)|, x \in [-1, 1]\}$. In the case where g is linear $|r_t| \leq \frac{1}{\lambda}$.

This suggest that the (normalized) market depth λ should not be too large in order to allow for realistic range of daily returns. We have chosen here $5 \leq \lambda \leq 20$ which allows a (maximal) range of daily returns between 5% and 20%. Note that this is a maximal range and, in fact, we will see that this maximum is not attained in a typical sample path. In practice, varying λ within this range does not affect the qualitative properties of the return process.

As noted above, $1/s$ represents the average number of periods an agent takes to update her views on market volatility: as noted in Section 4, unless the market is entirely composed of professional day traders following the market on a daily basis (for which $s = 1$), it is more plausible that $s \ll 1$. We have chosen here s within the range $s = 10^{-3} - 10^{-2}$.

The amplitude D of the input noise can be chosen as to reproduce a realistic range of values for the (annualized) volatility, such as measured by the moving average indicator $\hat{\sigma}$ defined above. Intuitively, the volatility of returns is increasing in D (although the relation turns out to be nonlinear) and this constraint leads us to choosing D in the range $10^{-3} - 10^{-2}$.

Finally, we choose $N = 10^3 - 10^4$ in order to get a realistic number of investors in a market. When the number of agents is lowered, the distribution of returns becomes multimodal with 3 local maxima, one at zero, one positive maximum and a negative one (figure2). This can be interpreted as a disequilibrium regime: the market moves either one way or the other. Stanley et al. (2003) have observed empirically a similar behavior.

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. These results discussed in the sequel are generic within this range of parameters, which simply make a comparison with daily returns possible.

3.3 Simulation results.

Using the parameter ranges above, we have performed an extensive simulation study of price behavior in this model. 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, the figures represent also the histogram of returns both in linear and log scales, the ACF of returns C_r, the ACF of absolute returns $C_{|r|}$.
More interestingly, we note that all the return series obtained possess some regularities which match some empirical properties outlined in Section 1.1:

•Excess volatility: the sample standard deviation of returns can be much larger than the standard deviation of the input noise representing news arrivals $\hat{\sigma}(t) \gg D$.

•Mean-reverting volatility: the market price fluctuates endlessly and displays "stochastic volatility": the volatility, as measured by the moving average estimator $\hat{\sigma}(t)$, does neither go to zero nor to infinity and displays a mean-reverting behavior. This behavior is attested by many empirical studies and GARCH models (e.g., Engel, 1995) on one hand and stochastic volatility models on the other hand aim at reproducing this mean-reverting stochastic behavior of volatility.

•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 1-3, the tail exhibit an approximately exponential decay, as observed in various studies of daily returns (e.g., Engle et al., 1993). Note that $\kappa = 6$ for a (two-sided) exponential distribution.

•The returns are uncorrelated: the sample autocorrelation function of the return exhibits an insignificant value (very similar to that of asset returns) at all lags, indicate absence of linear serial dependence in the returns.

•Volatility clustering: the autocorrelation function of absolute returns remains positive, and significantly above the autocorrelation of the returns, over many time lags, corresponding to persistence of the amplitude of returns a time scale ranging from a few weeks to several months. This is an indication of nonlinear dependence in the returns.

We will now attempt to understand the origin of these properties by analyzing the model.

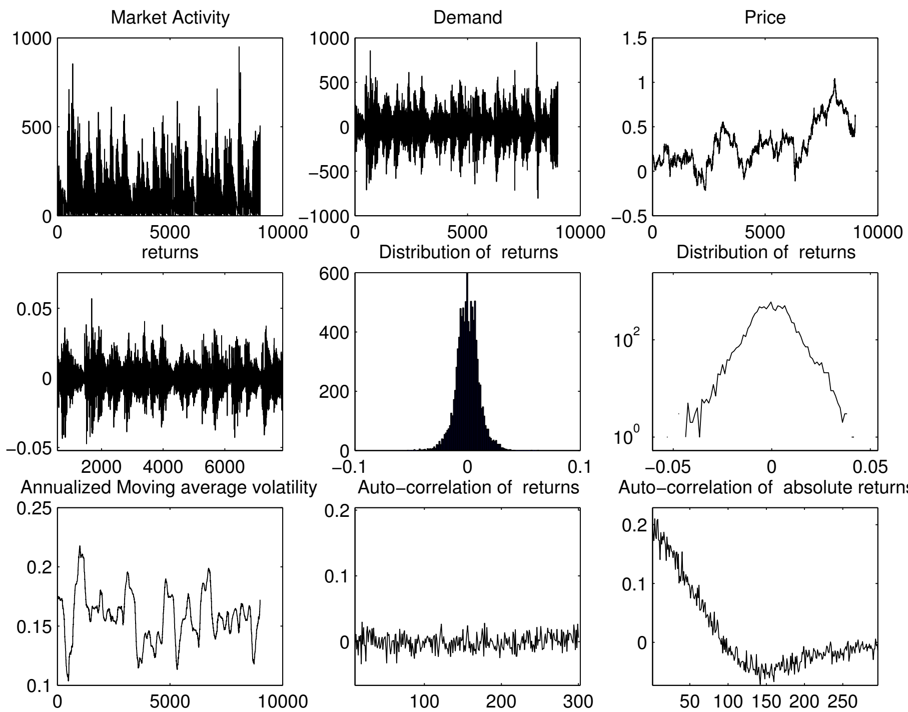

Figure 1: Numerical simulation of the model with updating frequency $s = 0.015$ (updating period: 67 "days"), $N = 1500$ agents, $\lambda = 10$, $D = 0.001$.

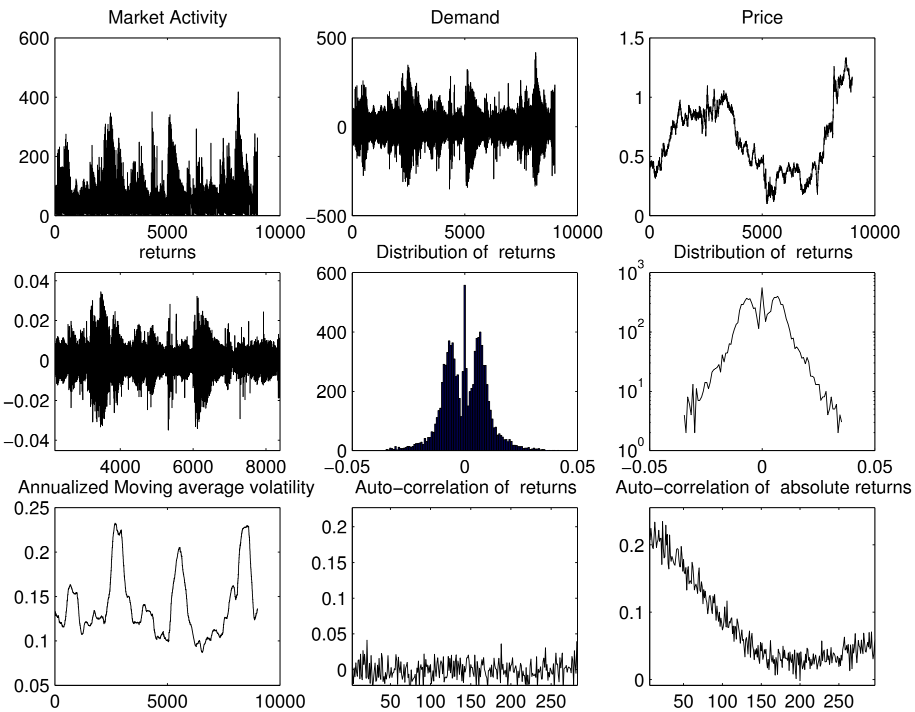

Figure 2: Numerical simulations of the model with updating frequency $s = 0.01$ (updating period: 100 days), $N = 1000$, $\lambda = 10$, $D = 0.001$.

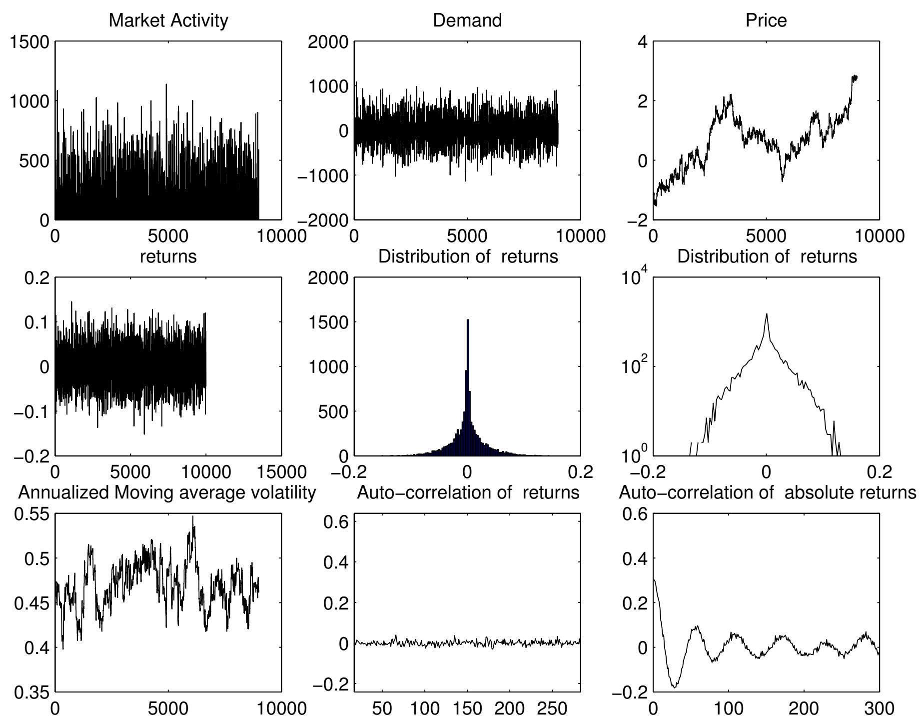

Figure 3: Numerical simulations of the model with updating frequency $s = 0.1$ (updating period: 10 days), $N = 1500$, $\lambda = 5$, $D = 0.001$.

4 Some limiting cases.

Here we study some limiting cases of the model, which allow to deduce some simple properties for the returns.

4.1 Feedback without heterogeneity: $s = 1$.

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 absolute return of the last period:

$$\theta_i(t) = |r_{t-1}|. \tag{10}$$

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 $|r_{t-1}|$:
$r_t = f(|r_{t-1}|, \epsilon_t) = g(1)1_{\epsilon_t > |r_{t-1}|} + g(-1)1_{\epsilon_t < -|r_{t-1}|}$,
a dependence structure typical of ARCH models, 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:

$$r_t \in \{0, g(1), g(-1)\}. \tag{11}$$

Due to the special threshold structure of the model, the autocorrelation of the absolute returns is in fact negative at lag 1: indeed, when the return take the value 0, the thresholds take the value 0 and at the next time step agents will act with a probability 1 so the return will take a non zero value $\pm 1/\lambda$ and the thresholds will be updated at $1/\lambda$. For D small enough compared to $1/\lambda$, the agents will not act with a probability 1 and again the returns will take the value 0.

Of course, this behavior is extreme since it implies that the agents have identical trading strategies and all agents are trading at each period, saturating market activity: it does not allow for market inertia. It also generates a trinomial distribution of returns which is not realistic. Simulation studies show that a similar behavior persists for $1 - s \ll 1$, leading to tri-modal distributions. This confirms our intuition that the updating probability s, which reflects the proportion of agents updating their choices at a given period, should be chosen small in order to guarantee the heterogeneity of the population.

4.2 Heterogeneity without feedback: s=0.

In the case where $s = 0$, no updating takes places: the trading strategies, given by the thresholds θ_i, are unaffected by the price behavior and the *feedback* effect is not present anymore. 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:

$$r_t = g(\frac{1}{N}\sum_{i=1}^{N} 1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}) = F(\epsilon_t). \tag{12}$$

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 (12), whose properties depend on the (initial) distribution of thresholds $(\theta_i, i = 1..N)$. The log-price then follows a (non-Gaussian) random walk and the model does not exhibit volatility clustering.

5 Behavior 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. Let us now turn to the general case.

5.1 General remarks.

Define $a_k = g(k/N)$ for $k = -N..0..N$ and consider the finite sets

$$E = \{a_k, k = -N..0..N\}, E_+ = \{g(k/N), k = 0..N\} = \{0, a_1, ..a_N\} \tag{13}$$

The returns take their value in E while the thresholds $\theta_i(t)$ take their values in E_+. Let us start by noting that the law of the thresholds $(\theta_i(t), i = 1..N)$ only depends on their values at t-1. We have with a probability 1-s:

$$\theta_i(t+1) = \theta_i(t) \tag{14}$$

and with probability s

$$\theta_i(t+1) = |r_t| = |g(\frac{1}{N}\sum_i [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}])|. \tag{15}$$

In the extreme cases where s=0 or s=1, we observed that the return r_t was in fact a Markov chain. This is not true in general: as seen in the above relations, the return r_t not only depends on r_{t-1} but also on the states of the agents $\theta_i(t)$. However it is readily observed from the above that:

Property 1 $[\theta_i(t), i = 1..N]$ *is a Markov chain in* E_+^N.

More interestingly, given that agents are indistinguishable and only the empirical distribution of threshold values affects the returns, defining $N_k(t)$ as the (random) number of agents with a threshold smaller than the k-th value $N_k(t) = \sum_{i=1}^{N} 1_{[0,a_k[}(\theta_i(t))$ one can actually show that the Markovian dynamics can be entirely described by $(N_k(t), k = 0..N)$:

Property 2 (Evolution of heterogeneity) $(N_k(t), k = 0..N)_{t=0,1..}$ *evolves as a Markov chain in* $\{0, ..., N\}^N$.

Note that $N(t) = (N_k(t), k = 0..N)$ 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 $\{0, ...N\}$, which is characteristic of disordered systems. In this case, even if we start from a deterministic set of values for the initial thresholds (even identical ones), the population distribution will evolve. By contrast with some models of disorder which have been used as analogies for systems of economic agents (e.g., Challet et al., 2001), here the disorder is endogeneous and is generated by the random updating mechanism.

5.2 Excess volatility

By "excess volatility" one refers to the observation that the level of variability in market prices is much higher than what can be expected based on the variability of fundamental economic variables and is unexplained by news arrivals. 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,

$$r_t = g(\frac{1}{N}\sum_{i=1}^{N}[1_{\epsilon_t>\theta_i} - 1_{\epsilon_t<-\theta_i}]) \tag{16}$$

the relation between these two quantities is far from being an equality: even after conditioning on the current states of agents $\theta_i(t), i = 1..N$, Eq. (16) 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 practice one can compare D with a commonly used notion of "volatility", a moving average estimator of standard deviation of returns, as decribed in Section 3. 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. This is a generic phenomenon also observed in other simulations: the order of magnitude of the volatility of returns may be totally different from that of the input noise. Here the mechanism for generating this excess volatility is identified as the threshold behavior of individual agents along with the heterogeneity of their behavioral rules.

5.3 Dependence properties of returns.

The following equations link the returns r_t, the thresholds $[\theta_i(t), i = 1..N]$ and the input noise ϵ_t:

$$Z_t = \sum_i \phi_i(t) = \sum_i [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}]. \tag{17}$$

$$r_t = g(Z_t/N) = g(\frac{1}{N} \sum_i [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}]). \tag{18}$$

From these relations one deduces the following properties:

Property 3 (Uncorrelated returns) *Assume that g is an odd function. Asset returns $(r_t)_{t \geq 0}$ are uncorrelated: $cov(r_t, r_{t+1}) = 0$.*

Property 4 (Volatility clustering) *Amplitudes of consecutive returns are positively correlated: $cov(|r_t|, |r_{t+1}|) > 0$.*

Of course, the volatility clustering property is observed to hold well beyond the first lag in the simulations shown above. In fact, defining the cluster length τ_c as the first lag for which the autocorrelation of absolute returns becomes zero, one can inquire into the dependence of this cluster length with respect to the updating frequency s. Figure 4 shows that the duration τ_c of volatility clusters increases with the average updating time 1/s in a monotone fashion and as a first approximation $\tau \simeq 1/s$. This interpretation is interesting since it related an observable quantity, τ_c, to the parameter s which describes the updating behavior of the agents. In most markets, the length of volatility cluster is roughly of the order of months, indicating that the range $s \sim 10^{-2}$ is in fact quite consistent with this behavior.

5.4 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 refered to as *investor inertia*, was used by Cont & Bouchaud (2000) as a key ingredient for generating long-memory properties in returns. However, while in these models investor inertia is exogeneously assumed, in the present case it is an outcome of the model: although all agents have the *possibility* of trading at each period, we observe that only a fraction of agents actually submit orders.
Starting from an initial holding of $p_i(0)$, the quantity of asset held by agent i is given by $p_i(t) = \sum_{\tau=0}^{t} \phi_i(\tau)$. Figure 5 displays the evolution of the portofolio $p_i(t)$ of a typical agent. The phenomenon of investor inertia is readily observed in this figure: short periods of activity (trading) are separated by long periods

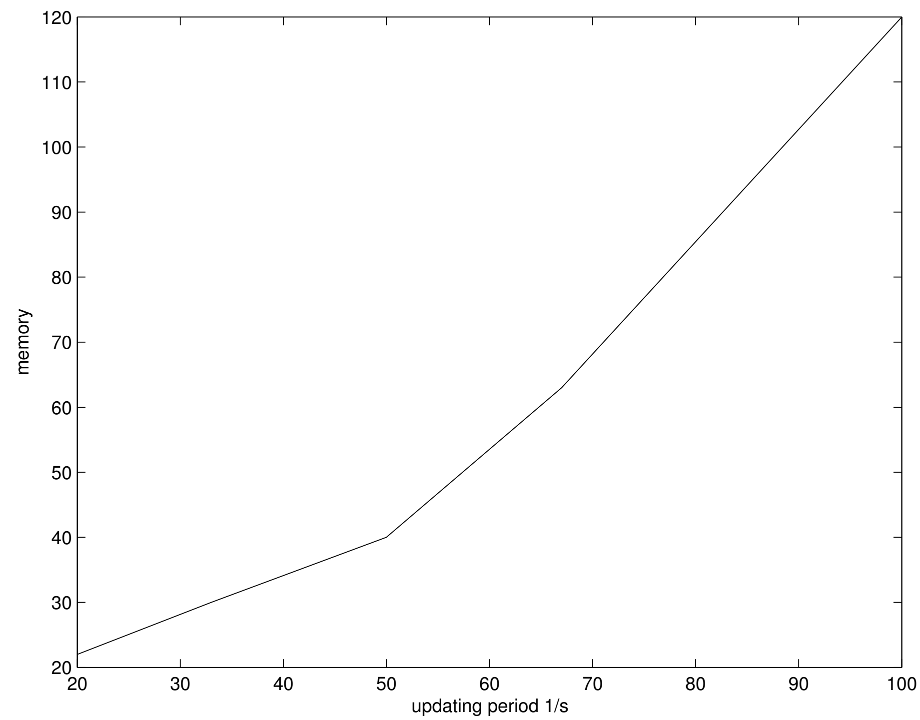

Figure 4: Time scale τ_c over which absolute returns remain positively correlated, as a function of the updating period 1/s.

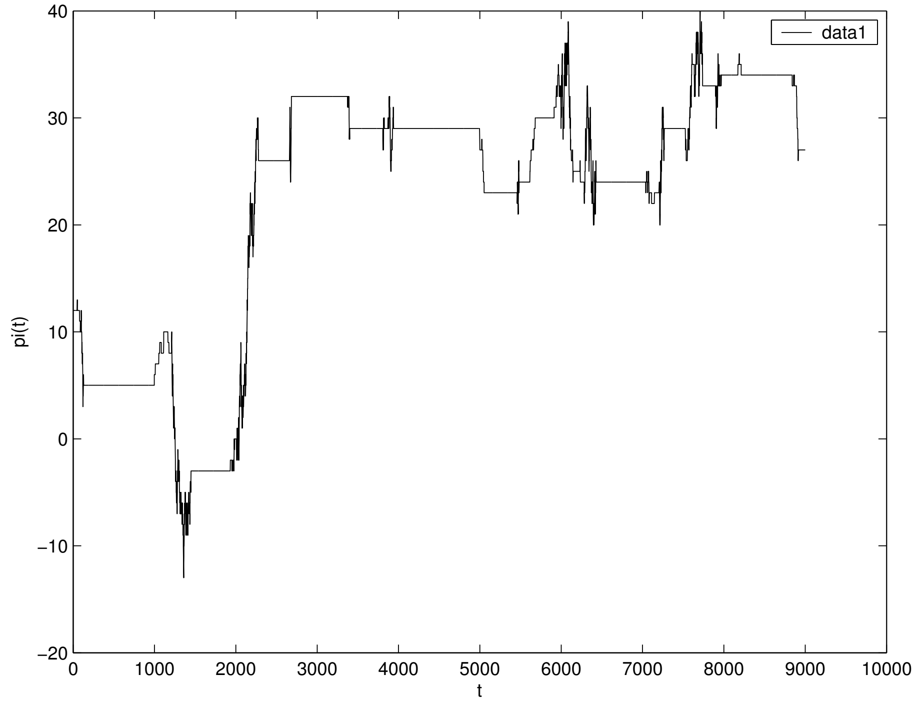

Figure 5: Evolution of the portofolio of a typical agent, in the market simulation displayed in figure1.

of inactivity, where the portofolio remains constant. Moreover, this "inertia" increases in periods of high volatility: agents update their thresholds to high values of absolute returns and such high values of thresholds makes them less reactive to the arrival of information and decreases order flow. While there is no explicit modeling of preferences or risk aversion in our simplified model, this effect is similar to the behavior of risk-averse agent maximizing the expectation of a concave utility.
Note that we are studying daily fluctuations and thus neglecting the full details of the mechanism of transaction and the problem of optimization of each agent's portofolio. Indeed, we describe one type of agent assuming that they are the main responsible of the price fluctuations of a single asset, but we do not describe the whole market. However, it would be interesting to check the effects of adding constraints on the agent's behavior (budget constraint for example), and this will be investigated in an other work.

5.5 Mean-reverting volatility.

Many microstructure models-especially those with learning or evolution- when observed over large time intervals, converge to an equilibrium where prices and other aggregate quantities cease to fluctuate randomly. Of course, this is not observed in financial markets: prices fluctuate endlessly and the volatility exhibits mean-reverting behavior.

It is easy to understand why our agent-based model will generically lead to endless price fluctuations and mean-reverting volatility. The mechanism is the following. Suppose we are in a period of "low-volatility"; the amplitude $|r_t|$ of returns is small. Agent 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.

The persistence of a low or high volatility period depends on how frequently agents update their thresholds, thus reacting to marklet activity. If s is the proportion of agents updating their thresholds at each period, the duration of such periods is of order 1/s.

The mechanism of this mean-reverting behavior can be understood in more detail in the case where the amplitude D of the input noise ϵ_t is small. Assume for simplicity that ϵ_t has bounded support and $P(|\epsilon_t| > g(1/N)) = 0$. Let us note $N_0(t)$ the number of agents with threshold zero at time t. A fraction of them will update their threshold to a non zero value, $sN_0(t)$ is the most probable fraction of updating agents so :

$$E[N_0(t+1)] = (1-s)N_0(t). \tag{19}$$

until we have zero agents with threshold zero, then again we expect sN agents to get their thresholds at zero. Because $r_t = g(\pm\frac{N_0(t)}{N})$, and with g linear, if $|r_t| \neq 0$

$$E[|r_{t+1}|||r_t|] = (1-s)|r_t| \tag{20}$$

and from this we get :

$$E[|r_{t+j}|||r_t|] = (1-s)^j|r_t|. \tag{21}$$

until we have no agents with threshold zero.

In figure 6, we have a simulation for $D = 10^{-5}$. It confirms the dynamics described above : bursts of activity around the value sN and then exponential decays, the process is not periodic however the sequences of activity last around a period of value $-\frac{1}{ln(1-s)}$.

From this, one can guess why the autocorrelation function of the absolute returns is positive and decays on a period that is an increasing function of the updating period $\frac{1}{s}$ as observed in figure 7, however this is not obvious to prove.

One can write :

$$N_0(t+1) = \sum_{i=1}^{N_0(t)} 1_{\gamma_i \leq 1-s} + \delta(N_0(t)) \sum_{j=1}^{N} 1_{\gamma_j \leq s} \tag{22}$$

with γ_i and γ_j random variables chosen uniformly on $[0,1]$. From this we get : The two-point autocorrelation function of the absolute returns is given by :

$$\varrho(1) = \frac{cov(N_0(t+1)N_0(t))}{\sqrt{var(N_0(t+1))var(N_0(t))}} \tag{23}$$

From (22) we prove that the two-point autocorrelation function is positive :

$$\varrho(1) = (1-s)\sqrt{\frac{var(N_0(t))}{var(N_0(t+1))}} > 0. \tag{24}$$

Thus in the case of D small enough, volatility decays exponentially in time and increases through upward "jumps" of magnitude sN. This behavior is actually similar to that of a class of stochastic volatility models, introduced by Barndorff-Nielsen & Shepard (2001) and successfully used to describe various econometric properties of returns.

One can understand qualitatively the dynamics of heterogeneity and how this is linked to the stylized facts. When a majority of agents have a low value for their threshold, it is very probable to get a large fluctuation. Because only a small fraction do higher their threshold response when a large fluctuation occurs, it is still very probable to get a large fluctuation at the next time step. In other words, the slow feedback mechanism causes persistence in the fluctuations.

6 Conclusion.

We have presented a parsimonious agent-based model which is capable of reproducing the main empirical stylized facts observed in returns of financial assets. Our model is based on three main ingredients:

•Threshold behavior of agents.

•Heterogeneity of agent strategies, generated endogeneously through random asynchronous updating of thresholds.

•A feedback effect of recent price behavior on agent strategies.

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 ingredients suffice for reproducing several empirical stylized facts as heavy tails, absence of autocorrelation in returns and volatility clustering, with realistic values in the time scales involved. In

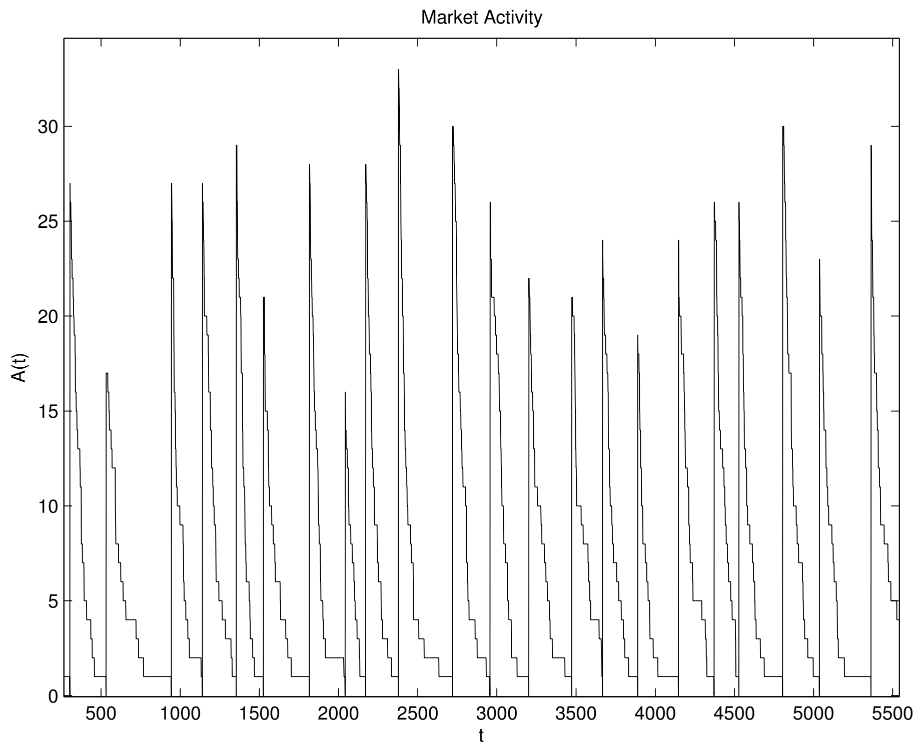

Figure 6: Numerical simulation of the model with $s = 0.015$ $N = 1500$ $\lambda = 0.1$ $D = 10^{-5}$

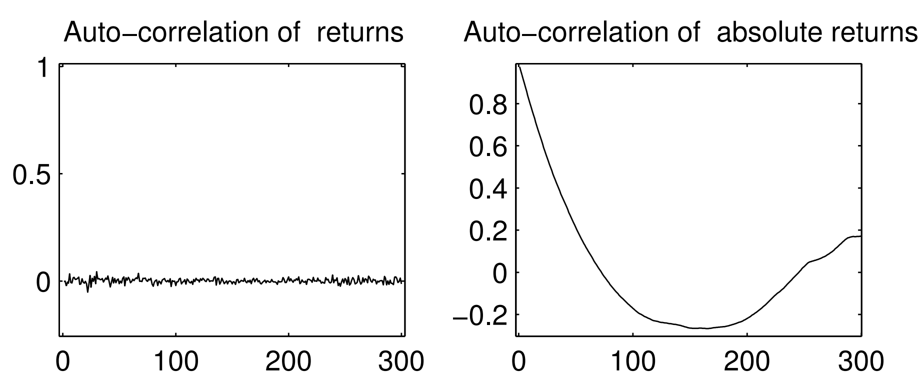

Figure 7: Numerical simulation of the model with $s = 0.015$ $N = 1500$ $\lambda = 0.1$ $D = 10^{-5}$

particular we do not introduce any exogeneous "fundamental" price or any direct interaction between agents, nor do we use different classes of traders (such as "chartists" or "fundamentalists"). These results question some conclusions previously drawn from simulation of agent-based models regarding the origins of stylized properties of asset returns and call for a closer, critical look at this issue through the study of a wider variety of agent-based market designs. We hope that the present work will contribute to this objective.

References

Arthur, W.B., Holland, J.H., 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 II, (Reading MA: Perseus Book), 15-44.

Banerjee, A., 1993. *The economics of rumors*, Review of Economic Studies 60, 309-327.

Barndorff-Nielsen, O.E., Shepard, N., 2001. *Non-Gaussian Ornstein-Uhlenbeck based models and some of their uses in financial econometrics*, Journal of the Royal Statistical Society B 63, 167-241.

Bayraktar, E., Horst, U., Sircar, R., 2003. *A limit theorem for financial markets with inert investors*, working paper, princeton University.

Blume, L., Easley, D., 1992. *Evolution and market behavior*, Journal of economic Theory 58, 9-40.

Brock, W.A., Hommes, C., 1998. *Heterogeneous beliefs and route to chaos in a simple asset pricing model*, Journal of Economic Dynamics and Control 22, 1235-1274.

Challet, D., Chessa, A., Marsili, M., Zhang, Y., 2001. *From minority games to real markets*, Quantitative Finance 1, 168-176.

Chiarella, C., He, X.-Z, 2002. *Heterogeneous Beliefs, Risk and Learning in a Simple Asset Pricing Model*, Computational Economics 19, 95-132.

Cont, R., 2001. *Empirical properties of asset returns: Stylised facts and statistical issues*, Quantitative Finance 1, 223-236.

Cont, R., Bouchaud, J.-P., 2000. *Herd Behavior and aggregate fluctuations in financial markets*, Macroeconomic dynamics 4, 170-196.

Cutler, D., Poterba, J., Summers, L., 1989. *What moves stock prices?*, Journal of Portofolio Management, 4-12.

Ding, Z., Granger, C., Engle, R., 1993. *A long memory property of stock market returns and a new model*, Journal of empirical finance 1, 83-106.

Engle, R., 1995. *ARCH: Selected Readings*, Oxford University press, Oxford.

Farmer, J.D., Joshi, S., 2002. *The price dynamics of common trading strategies*, Journal of Economic Behavior and Organization 49, 149-71.

Gaunersdorfer, A., 2000. *Endogeneous fluctuations in a simple asset pricing model with heterogeneous agents*, Journal of Economic Dynamics and Control 24, 799-831.

Granovetter, M., Soong, R., 1983. *Threshold models of diffusion and collective behavior*, Journal of mathematical sociology, 165-179.

Guillaume, D., Dacarogna, M., Davé, R., Müller, U., Olsen, R., 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.

Hommes, C.H., Gaunersdorfer, A., Wagener, F.O., 2003. *Bifuraction routes to volatility clustering under evolutionary learning*, technical report, CenDEF.

Hong, H., Kubik, J.D., Solomon, A., 2000. *Security analysts' career concerns and herding of earning forecasts*, RAND Journal of Economics 31, 121-144.

Kirman, A., Teyssiere, G., 2002. *Microeconomic models for long-memory in the volatility of financial time series*, Studies in nonlinear dynamics and econometrics 5, 281-302.

LeBaron, B., 2000. *Agent-based computational finance : suggested readings and early research*, Journal of Economic Dynamics and Control 24, 679-702.

LeBaron, B., 2001, *Evolution and time horizons in an agent-based stock market*, Macroeconomic Dynamics 5, 225-254.

Lux, T., Marchesi, M., 2000. *Volatility clustering in financial markets : a micro simulation of interacting agents*, International Journal of Theoretical and Applied Finance 3, 675-702.

Mandelbrot, B.B., 1963. *The variation of certain speculative prices*, Journal of Business 36, 394-419.

Orléan, A., 1995. *Bayesian interactions and collective dynamics of opinions*, Journal of Economic behavior and Organization 28, 257-74.

Shiller, R., 2000. *Irrational Exuberance*, Princeton University Press.

Stauffer, D., 2001. *Percolation models of financial market dynamics*, Advances in Complex Systems 4, 19-27.

Stauffer, D., de Oliveira, P.M.C., Bernardes, A.T., 1999. *Monte Carlo simulation of volatility clustering in market model with herding*, International Journal of theoretical Finance 2, 83-94.

Zovko, I., Farmer, J.D., 2002. *The power of patience: a behavioral regularity in limit order placement*, Quantitative Finance 2, 387-392.

Plerou, V., Gopikrishnan, P., Stanley, H.E., 2003. *Two-phase behavior of financial markets*, Nature 421, 130.