
Volatility Clustering in Financial Markets: Empirical Facts and Agent–Based Models

Rama Cont*

Centre de Mathématiques Appliquées, Ecole Polytechnique,
F-91128 Palaiseau, France. Rama.Cont@polytechnique.fr

Summary. Time series of financial asset returns often exhibit the *volatility clustering* property: large changes in prices tend to cluster together, resulting in persistence of the amplitudes of price changes. After recalling various methods for quantifying and modeling this phenomenon, we discuss several economic mechanisms which have been proposed to explain the origin of this volatility clustering in terms of behavior of market participants and the news arrival process. A common feature of these models seems to be a switching between low and high activity regimes with heavy-tailed durations of regimes. Finally, we discuss a simple agent-based model which links such variations in market activity to threshold behavior of market participants and suggests a link between volatility clustering and investor inertia.

1 Introduction

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 (Ding *et al.*, 1993, Guillaume *et al.*, 1997, Pagan, 1996, 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 occurrence of large (negative or positive) returns is not always explainable by the arrival of new information on the market (Cutler *et al.*, 1989).
- **Heavy tails:** the (unconditional) distribution of returns displays a heavy tail with positive excess kurtosis.

* The author thanks Alan Kirman and Gilles Teyssière for their infinite patience and participants in the CNRS Summer School on Complex Systems in the Social Sciences (ENS Lyon, 2004) for their stimulating feedback. The last section of this paper is based on joint work with F. Ghoulmie and J.P. Nadal.

- **Absence of autocorrelations in returns:** (linear) autocorrelations of asset returns are often insignificant, except for very small intraday time scales ($\simeq 20$ minutes) where 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 $|r_t|$ or their squares display a positive, significant and slowly decaying autocorrelation function: $\text{corr}(|r_t|, |r_{t+\tau}|) > 0$ for τ ranging from a few minutes to a several weeks.
- **Volume/volatility correlation:** trading volume is positively correlated with market volatility. Moreover, trading volume and volatility show the same type of “long memory” behavior (Lobato and Velasco, 2000).

Among these properties, the phenomenon of volatility clustering has intrigued many researchers and oriented in a major way the development of stochastic models in finance –GARCH models and stochastic volatility models are intended primarily to model this phenomenon. Also, it has inspired much debate as to whether there is long-range dependence in volatility. We review some of these issues in Section 2. As noted by the participants of this econometric debate (Willinger *et al.*, 1999, Mikosch and Stărică, 2003), statistical analysis alone is not likely to provide a definite answer for the presence or absence of long-range dependence phenomenon in stock returns or volatility, unless economic mechanisms are proposed to understand the origin of such phenomena.

Some insights into these economic mechanisms are given by agent-based models of financial markets. Agent-based market models attempt to explain the origin of the observed behavior of market prices in terms of simple, stylized, behavioral rules of market participants, (Chiarella *et al.*, 2003, Lux, 1998, Lux and Marchesi, 2000, LeBaron 2001a): 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. We review some of these approached in Section 3 and discuss how they lead to volatility clustering.

Most of these agent-based models are complex in structure and have been studied using Monte Carlo simulations. As noted also by LeBaron (2000), due to the complexity of such models 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. In Section 4 we present an agent-based model capable of generating time series of asset returns with properties similar to the stylized facts above, but which is simple enough in structure so the origins of volatility clustering can be traced back to agents behavior. This model points to a link between investor inertia and volatility clustering and provide an economic explanation for the

switching mechanism proposed in the econometrics literature as an origin of volatility clustering.

2 Volatility Clustering in Financial Time Series

Denote by S_t the price of a financial asset — a stock, an exchange rate or a market index — and $X_t = \ln S_t$ its logarithm. Given a *time scale* Δ , the log return at scale Δ is defined as:

$$r_t = X_{t+\Delta} - X_t = \ln\left(\frac{S_{t+\Delta}}{S_t}\right). \quad (1)$$

Δ may vary between a minute (or even seconds) for tick data to several days. Observations are sampled at discrete times $t_n = n\Delta$. Time lags will be denoted by the Greek letter τ ; typically, τ will be a multiple of Δ in estimations. For example, if $\Delta = 1$ day, $\text{corr}[r_{t+\tau}, r_t]$ denotes the correlation between the daily return at period t and the daily return τ periods later.

2.1 Empirical Behavior of Autocorrelation Functions

A typical display of daily log-returns is shown in figure 1: the volatility clustering feature is seen graphically from the presence of sustained periods of high or low volatility. As noted above, the autocorrelation of returns is typically insignificant at lags between a few minutes and a month. An example is shown in figure 2 (left). This “spectral whiteness” of returns can be attributed to the activity of arbitrageurs who exploit linear correlations in returns via trend following strategies, see Mandelbrot (1971). By contrast, the autocorrelation function of absolute returns remains positive over lags of several weeks and decays slowly to zero: figure 2 (right) shows this decay for SLM stock (NYSE). This observation is remarkably stable across asset classes and time periods and is regarded as a typical manifestation of volatility clustering (Bollerslev *et al.*, 1992, Ding *et al.*, 1993, Cont *et al.*, 1997, Guillaume *et al.*, 1997). Similar behavior is observed for the autocorrelation of squared returns, see Bollerslev *et al.* (1992), and more generally for $|r_t|^\alpha$ (Ding *et al.*, 1993, Ding and Granger, 1996, Cont *et al.*, 1997), but it seems to be most significant for $\alpha = 1$, i.e. absolute returns (Ding *et al.*, 1993).

GARCH models (Bollerslev *et al.*, 1992, Engle, 1995) were among the first models to take into account the volatility clustering phenomenon. In a GARCH(1,1) model the (squared) volatility depends on last periods volatility:

$$r_t = \sigma_t \varepsilon_t, \quad \sigma_t^2 = a_0 + a\sigma_{t-1}^2 + b\varepsilon_t^2, \quad 0 < a + b < 1, \quad (2)$$

leading to positive autocorrelation in the volatility process σ_t , with a rate of decay governed by $a + b$: the closer $a + b$ is to 1, the slower the decay of the autocorrelation of σ_t . The constraint $a + b < 1$ allows for the existence of a

BMW stock daily returns




Fig. 1. Large changes cluster together: BMW daily log-returns. $\Delta = 1$ day

stationary solution, while the upper limit $a + b = 1$ corresponds to the case of an integrated process. Estimations of GARCH(1,1) on stock and index returns usually yield $a + b$ very close to 1 (Bollerslev *et al.*, 1992). For this reason the volatility clustering phenomenon is sometimes called a “GARCH effect”; one should keep in mind however that volatility clustering is a “non-parametric” property and is not intrinsically linked to a GARCH specification.

While GARCH models give rise to exponential decay in autocorrelations of absolute or squared returns, the empirical autocorrelations are similar to a power law; see Cont *et al.* (1997), Guillaume *et al.* (1997):

$$C_{|r|}(\tau) = \text{corr}(|r_t|, |r_{t+\tau}|) \simeq \frac{c}{\tau^\beta},$$

with an exponent $\beta \leq 0.5$, (Cont *et al.*, 1997, Breidt *et al.*, 1998), which suggests the presence of “long-range” dependence in amplitudes of returns, discussed below.




Fig. 2. SLM stock, NYSE, $\Delta = 5$ minutes. Left: autocorrelation function of log-returns. Right: autocorrelation of absolute log-returns

2.2 Long Range Dependence

Let us recall briefly the commonly used definitions of long range dependence, based on the autocorrelation function of a process:

Definition 1 (Long range dependence). A stationary process Y_t (with finite variance) is said to have long range dependence if its autocorrelation function $C(\tau) = \text{corr}(Y_t, Y_{t+\tau})$ decays as a power of the lag τ :

$$C(\tau) = \text{corr}(Y_t, Y_{t+\tau}) \underset{\tau \rightarrow \infty}{\sim} \frac{L(\tau)}{\tau^{1-2d}}, \quad 0 < d < \frac{1}{2}, \tag{3}$$

where L is slowly varying at infinity, i.e. verifies $\forall a > 0, \frac{L(at)}{L(t)} \rightarrow 1$ as $t \rightarrow \infty$.

By contrast, one speaks of “short range dependence” if the autocorrelation function decreases at a geometric rate:

$$\exists K > 0, c \in]0, 1[, \quad |C(\tau)| \leq Kc^\tau. \tag{4}$$

Obviously, (3) and (4) are not the only possibilities for the behavior of the autocorrelation function at large lags: there are many other possible decays rates, intermediate between a power decay and a geometric decay. However, it is noteworthy that in all stochastic models used in the financial modeling literature, the behavior of returns and their absolute values fall within one of the two categories.

The long range dependence property (3) hinges upon the behavior of the autocorrelation function at *large* lags, a quantity which may be difficult to estimate empirically, see Beran (1994). For this reason, models with long-range dependence are often formulated in terms of self-similar processes, which allow

to extrapolate across time scales and deduce long time behavior from short time behavior, which is more readily observed. A stochastic process $(X_t)_{t \geq 0}$ is said to be self-similar if there exists $H > 0$ such that for any scaling factor $c > 0$, the processes $(X_{ct})_{t \geq 0}$ and $(c^H X_t)_{t \geq 0}$ have the same law:

$$(X_{ct})_{t \geq 0} \stackrel{d}{=} (c^H X_t)_{t \geq 0}. \quad (5)$$

H is called the self-similarity exponent of the process X . Note that a self-similar process cannot be stationary, so the above definition of long-range dependence cannot hold for a self-similar process, but eventually for its increments (if they are stationary). The typical example of self-similar process whose increments exhibit long range dependence is fractional Brownian motion (Mandelbrot and Van Ness, 1968).

But self-similarity does not imply long-range dependence in any way: α -stable Lévy processes provide examples of self-similar processes with *independent* increments. Nor is self-similarity implied by long range dependence: Cheridito (2004) gives several examples of Gaussian semimartingales with the same long range dependence features as fractional Brownian noise but with no self-similarity (thus very different “short range” properties and sample path behavior). The example of fractional Brownian motion is thus misleading in this regard, since it conveys the idea that these two properties are associated. When testing for long range dependence in a model based on fractional Brownian motion, we thus test the joint hypothesis of self-similarity *and* long-range dependence and strict self-similarity is not observed to hold in asset returns (Cont *et al.*, 1997, Cont, 2001).

A fallacy often encountered in the literature is that long range dependence in returns is incompatible with absence of (continuous-time) arbitrage. Again, the origin of this idea can be traced back to models based on fractional Brownian motion: since fractional Brownian motion is not a semimartingale, a model in which the (log)-price are described by a fractional Brownian motion is not arbitrage-free (in the continuous-time sense); see Rogers (1997). This result (and the fact that fractional Brownian motions fails to be a semimartingale) crucially depends on the *local* behavior of its sample paths, not on its long range dependence property. Cheridito (2004) gives several examples of Gaussian processes with the same long range dependence features as fractional Brownian motion, but which are semimartingales and lead to arbitrage-free models.

2.3 Dependence in Stock Returns

The volatility clustering feature indicates that asset returns are not independent across time; on the other hand the absence of linear autocorrelation shows that their dependence is nonlinear. Whether this dependence is “short range” or “long range” has been the object of many empirical studies.

The idea that stock returns could exhibit long range dependence was first suggested by Mandelbrot (1971) and subsequently observed in many empirical studies using R/S analysis (Mandelbrot and Taqqu, 1979). Such tests have been criticized by Lo (1991) who pointed out that, after accounting for short range dependence, they might yield a different result and proposed a modified test statistic. Lo's statistic highly depends on the way "short range" dependence is accounted for and shows a bias towards rejecting long range dependence, see Teverovsky *et al.* (1999). The final empirical conclusions are therefore less clear (Willinger *et al.*, 1999).

However, the absence of long range dependence in returns may be compatible with its presence in absolute returns or "volatility". As noted by Heyde (2002), one should distinguish long range dependence in signs of increments, when $\text{sign}(r_t)$ verifies (3), from long range dependence in amplitudes, when $|r_t|$ verifies (3). Asset returns do not seem to possess long range dependence in signs (Heyde, 2002). Many authors have thus suggested models, such as FIGARCH (Baillie *et al.*, 1996), in which returns have no autocorrelation but their amplitudes have long range dependence; see Doukhan *et al.* (2003).

It has been argued (LeBaron, 2001b, Barndorff-Nielsen and Shephard, 2001) that the decay of $C_{|r_t|}(\tau)$ can also be reproduced by a superposition of several exponentials, indicating that the dependence is characterized by multiple time scales. In fact, an operational definition of long range dependence is that the time scale of dependence in a sample of length T is found to be of the order of T : dependence extends over the whole sample. Interestingly, the largest time scale in LeBaron (2001b) is found to be of the order of ... the sample size, a prediction which would be compatible with long-range dependence!

Many of these studies test for long range dependence in returns, volatility, etc. by examining sample autocorrelations, Hurst exponents etc. but if time series of asset returns indeed possess the two features of heavy tails *and* long range dependence, then many of the standard estimation procedures for these quantities may fail to work (Resnick, 1998). For example, sample autocorrelation functions may fail to be consistent estimators of the true autocorrelation of returns in the price generating process: Resnick *et al.* (1999) give examples of such processes where sample autocorrelations converge to *random* values as sample size grows! Also, in cases where the sample ACF is consistent, its estimation error can have a heavy-tailed asymptotic distribution, leading to large errors. The situation is even worse for autocorrelations of squared returns, see Mikosch and Střaričá (2000). Thus, one must be cautious in identifying behavior of *sample* autocorrelation with the autocorrelations of the return process.

Slow decay of sample autocorrelation functions may possibly arise from other mechanism than long-range dependence. For example, note that non-stationarity of the returns may also generate spurious effects which can be mistaken for long-range dependence in the volatility. However, we will not go to the extreme of suggesting, as in Mikosch and Střaričá (2003), that the slow

decay of sample autocorrelations of absolute returns is a pure artefact due to non-stationarity. “Non-stationarity” does not suggest a modeling approach and it seems highly unlikely that unstructured non-stationarity would lead to such a robust, stylized behavior for the sample autocorrelations of absolute returns, stable across asset classes and time periods. The robustness of these empirical facts call for an explanation, which “non-stationarity” does not provide. Of course, these mechanisms are not mutually exclusive: a recent study by Granger and Hyung (2004) illustrates the interplay of these two effects by combining an underlying long memory process with occasional structural breaks.

Independently of the econometric debate on the “true nature” of the return generating process, one can take into account such empirical observations without pinpointing a specific stochastic model by testing for similar behavior of sample autocorrelations in agent-based models (described below), and using sample autocorrelations for indirect inference (Gourieroux *et al.*, 1993), of the parameters of such models.

3 Mechanisms for Volatility Clustering

While GARCH, FIGARCH and stochastic volatility models propose statistical constructions which mimic volatility clustering in financial time series, they do not provide any economic explanation for it. We discuss here possible mechanisms which have been proposed for the origin of volatility clustering.

3.1 Heterogeneous Arrival Rates of Information

Heterogeneity in agent’s time scale has been considered as a possible origin for various stylized facts (Guillaume *et al.*, 1997). Long term investors naturally focus on long-term behavior of prices, whereas traders aim to exploit short-term fluctuations.

Granger (1980) suggested that long memory in economic time series can be due to the aggregation of a cross section of time series with different persistence levels. This argument was proposed by Andersen and Bollerslev (1997) as a possible explanation for volatility clustering in terms of aggregation of different information flows.

The effects of the diversity in time horizons on price dynamics have also been studied by Lebaron (2001a) in an artificial stock market, showing 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.

3.2 Evolutionary Models

Several studies have considered modeling financial markets by analogy with ecological systems where various trading strategies co-exist and evolve via

a “natural selection” mechanism, according to their relative profitability, (Arthur *et al.*, 1997, Arifovic and Gencay, 2000, LeBaron *et al.*, 1999, LeBaron, 2001a). The idea of these models, the prototype of which is the Santa Fe artificial stock market, (Arthur *et al.*, 1997, LeBaron *et al.*, 1999), is that a financial market can be viewed as a population of agents, identified by their (set of) decision rules. A decision rule is defined as a mapping from an agents information set (price history, trading volume, other economic indicators) to the set of actions (buy, sell, no trade). The evolution of agents decision rule is often modeled using a genetic algorithm (Holland, 1992). The specification and simulation of such evolutionary models can be quite involved and specialized simulation platforms have been developed to allow the user to specify variants of agents strategies and evolution rules. Other evolutionary models represent the evolution by a deterministic dynamical system which, through the complex price dynamics it generate, are able to mimick some “statistical” properties of the returns process, including volatility clustering; see Hommes *et al.* (2003).

Though the Santa Fe market model is capable of qualitatively replicating some of the stylized facts (LeBaron *et al.*, 1999) precise comparisons with empirical observations are still lacking. Indeed, given the large number of parameters, it is not possible to calibrate the parameters in order to interpret the time periods in the simulations as “days” or “minutes” etc. thereby leading to a lack of reference for empirical comparisons.

More importantly, the competition between numerous strategies in such complex simulation models does not allow to pinpoint a single mechanism as being responsible for volatility clustering or other stylized properties. Models in which a dominant mechanism is at work are more helpful in this respect; we will now discuss some instances of such models.

3.3 Behavioral Switching

The economic literature contains examples where switching of economic agents between two behavioral patterns leads to large aggregate fluctuations, Kirman (1993): in the context of financial markets, these behavioral patterns can be seen as trading rules and the resulting aggregate fluctuations as large movements in the market price i.e. heavy tails in returns. Recently, models based on this idea have also been shown to generate volatility clustering (Kirman and Teyssière, 2002, Lux and Marchesi, 2000).

Lux and Marchesi (2000) study an agent-based model in which heavy tails of asset returns and volatility clustering arise from behavioral switching of market participants between fundamentalist and chartist behavior. Fundamentalists expect that the price follows the fundamental value in the long run. Noise traders try to identify price trends, which results in a tendency to herding. Agents are allowed to switch between these two behaviors according to the performance of the various strategies. Noise traders evaluate their

performance according to realized gains, whereas for the fundamentalists, performance is measured according to the difference between the price and the fundamental value, which represents the anticipated gain of a “convergence trade”. This decision-making process is driven by an exogenous fundamental value, which follows a Gaussian random walk. Price changes are brought about by a market maker reacting to imbalances between demand and supply. Most of the time, a stable and efficient market results. However, its usual tranquil performance is interspersed by sudden transient phases of destabilization. An outbreak of volatility occurs if the fraction of agents using chartist techniques surpasses a certain threshold value, but such phases are quickly brought to an end by stabilizing tendencies. This behavioral switching is believed to be the cause of volatility clustering, long memory and heavy tails in the Lux-Marchesi (2000) model.

Kirman and Teyssière (2002) have proposed a variant of Kirman (1993) in which the proportion $\alpha(t)$ of fundamentalists in the market follows a Markov chain, of the type used in epidemiological models, describing herding of opinions. Simulation of this model exhibit autocorrelation patterns in absolute returns with a behavior similar to that described in Section 2.

3.4 The Role of Investor Inertia

As argued by Liu (2000), the presence of a Markovian regime switching mechanism in volatility can lead to volatility clustering, is not sufficient to generate long-range dependence in absolute returns. More important than the switching is the fact the time spent in each regime –the duration of regimes– should have a heavy-tailed distribution (Pourahmadi, 1988, Taqqu and Levy, 1986). By contrast with Markov switching, which leads to short range correlations, this mechanism has been called “renewal switching”.²

Bayraktar *et al.* (2003) study a model where an order flow with random, heavy-tailed, durations between trades leads to long range dependence in returns. When the durations τ_n of the inactivity periods have a distribution of the form $\mathbb{P}(\tau_n \geq t) = t^{-\alpha}L(t)$, conditions are given under which, in the limit of a large number of agents randomly submitting orders, the price process in this models converges to a process with Hurst exponent $H = (3 - \alpha)/2 > 1/2$. In this model the randomness (and the heavy tailed nature) of the durations between trades are both exogenous ingredients, chosen in a way that generates long range dependence in the returns. However, as noted above, empirical observations point to clustering and persistence in *volatility* rather than in returns so such a result does not seem to be consistent with the stylized facts.

By contrast, as noted above, regime switching in *volatility* with heavy-tailed durations could lead to volatility clustering. Although in the agent-based models discussed above, it may not be easy to speak of well-defined

² See the chapter by Giraitis, Leipus and Surgailis in this volume for a review on renewal switching models.

“regimes” of activity, but Giardina and Bouchaud (2003) argue that this is indeed the mechanism which generates volatility clustering in the Lux-Marchesi (2000) and other models discussed above. In these models, agents switch between strategies based on their relative performance; Giardina and Bouchaud argue that this (cumulative) relative performance index actually behaves in time like a random walk, so the switching times can be interpreted as times when the random walk crosses zero: the interval between successive zero-crossings is then known to be heavy-tailed, with a power-law decay of exponent $3/2$.

4 Volatility Clustering and Threshold Behavior

While switching between high and low volatility states is probably the mechanism leading to volatility clustering in many of the agent-based models discussed above, this explanation is not easy to trace back to the level of agent behavior, partly because the models described above contain various other ingredients whose contribution to the overall behavior is thus blurred. We now discuss a simple model (Cont *et al.*, 2004) reproducing several stylized empirical facts, where the origin of volatility clustering can be clearly traced back to investor inertia, caused by threshold response of investors to news arrivals.

4.1 An Agent-Based Model for Volatility Clustering

Our model describes a market where a single asset, whose price is denoted by S_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 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 $\phi_i(t) = -1$. We allow the value $\phi_i(t)$ to be zero; the agent is then inactive at period t . The inflow of public information is modeled 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 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 agents (subjective) view on volatility. The trading rule we study may be seen as a stylized example of threshold behavior: 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:

$$\phi_i(t) = 1_{\epsilon_t > \theta_i(t)} - 1_{\epsilon_t < -\theta_i(t)}. \quad (6)$$

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

$$r_t = \ln \frac{S_t}{S_{t-1}} = g\left(\frac{Z_t}{N}\right), \quad (7)$$

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 $g'(0) = \frac{1}{\lambda}$. Examples are a linear price impact $g(z) = z/\lambda$ or $g(z) = \arctan(z/\lambda)$, both having been used in various disequilibrium models.

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, q represents the fraction of agents updating their views at any period; $1/q$ 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, q 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 $|r_t| = \left| \ln \frac{S_t}{S_{t-1}} \right|$. 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:

$$\theta_i(t) = 1_{u_i(t) < s} |r_t| + 1_{u_i(t) \geq s} \theta_i(t-1). \quad (8)$$

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 (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. As noted above, the heterogeneity of time scales of intervention of agents is a feature believed to be important for generating persistence in volatility (Andersen and Bollerslev, 1997, Granger, 1980, LeBaron, 2000). The random nature of updating in this model is a parsimonious way to introduce heterogeneity in time scales 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:

1. agents receive a common signal $\epsilon(t) \sim N(0, D^2)$

2. each agent i compares the signal to her threshold $\theta_i(t)$
3. if $|\epsilon(t)| > \theta_i(t)$ the agent considers the signal as significant and generates an order $\phi_i(t)$ according to (6).
4. The market price is impacted by the excess demand and moves according to (7).
5. Each agent updates, with probability q , her threshold according to (8).

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: q 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 dynamics and properties similar to empirically observed properties of asset returns.

4.2 Simulation Results

In order for a direct comparison with empirical stylized facts to be meaningful, we compute sample moments as in the case of empirical data, by averaging over the (single) sample path. After simulating a sample path of the price S_t for $T = 10^4$ periods, we compute the time series of returns $r_t = \ln(S_t/S_{t-1}), t = 1, \dots, T$, 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 $|r_t| \leq \frac{1}{\lambda}$ 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 3 and 4 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 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 the introduction (Cont *et al.*, 2004):




Fig. 3. Numerical simulation of the model with updating frequency $q = 0.01$ (average updating period: 100 “days”) $N = 1000$ agents, $D = 0.001$ and $\lambda = 10$

1. Excess volatility: the sample standard deviation of returns can be much larger than the standard deviation of the input noise representing news arrivals $\hat{\sigma}(t) \gg D$.
2. Mean-reverting volatility: the market price fluctuates endlessly and the volatility, as measured by the moving average estimator $\hat{\sigma}(t)$, does neither to zero nor to infinity and displays a mean-reverting behavior.
3. 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 3–4, the tails exhibit an approximately exponential decay, as observed in various studies of daily returns (Ding *et al.*, 1993).
4. 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, indicating the absence of linear serial dependence in the returns.
5. Volatility clustering: the autocorrelation function of absolute returns remains significantly positive over many time lags, corresponding to persistence of the amplitude of returns a time scale $\simeq 1/q$.




Fig. 4. Numerical simulation of the model with updating frequency $q = 0.1$ (average updating period: 10 “days”) $N = 1500$ agents, $D = 0.001$ and $\lambda = 10$

4.3 Theoretical Analysis

Contrarily to some of the models discussed above, this model is simple enough to allow for a theoretical study of its qualitative studies (Cont *et al.*, 2004). Let us begin by examining two limiting cases:

1. **Feedback without heterogeneity:** In the case where $q = 1$, 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 (Engle, 1995), 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 behavior persists for $1 - q \ll 1$, leading to tri-modal distributions. This confirms our intuition that the updating probability q should be chosen small.

2. **Heterogeneity without feedback:** In the case where $q = 0$, no updating takes place: 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 :

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

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 (9), 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.

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 $0 < q \ll 1$: 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 $1 - q$ and

$$\theta_i(t + 1) = |r_t| = \left|g\left(\frac{1}{N} \sum_{i=1}^N [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}]\right)\right| \quad \text{with probability } q. \quad (9)$$

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 (Mézard *et al.*, 1984), 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 ,

$$r_t = g\left(\frac{\sum_{i=1}^N 1_{\epsilon_t > \theta_i(t)} - 1_{\epsilon_t < -\theta_i(t)}}{\lambda N}\right), \quad (10)$$

even after conditioning on the current states of agents $\theta_i(t)$, $i = 1, \dots, N$, equation (10) 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 3, $D = 10^{-3}$ 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

$$Z_t = \frac{1}{N} \sum_{i=1}^N \phi_i(t) = \frac{1}{N} \sum_{i=1}^N [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}], \quad (11)$$

$$r_t = g(Z_t) = g\left(\frac{1}{N} \sum_{i=1}^N [1_{\epsilon_t > \theta_i} - 1_{\epsilon_t < -\theta_i}]\right), \quad (12)$$

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 behavior 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 4.3 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 behavior of risk-averse agent.

- **Mean reversion and clustering of 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 behavior. 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




Fig. 5. Evolution of the portfolio of a typical agent, with long periods of inactivity punctuated by bursts of activity

the incoming signal: this increase in investor inertia will thus decrease the amplitude of returns. The mean reversion time in the volatility corresponds here to the time it takes for agents to adjust their thresholds to current market conditions, which is of order $\tau_c = 1/q$.

When the amplitude of the noise is small it can be shown (Cont *et al.*, 2004) that volatility decays exponentially in time and increases through upward “jumps”. This behavior is actually similar to that of a class of stochastic volatility models, introduced by Barndorff-Nielsen and Shephard (2001) and successfully used to describe various econometric properties of returns.

5 Conclusion

Volatility clustering is recognized as a stylized property present in most financial time series. Agent-based models seek to explain volatility clustering in terms of behavior of market participants, described in terms of simple rules. We have discussed several agent-based models capable of generating volatility clustering. A common feature of these models seems to be the “switching” of the market between periods of high and low activity, with long durations of periods. Models differ in the mechanism which leads to this switching at the level of agents.

While the econometric debate on the short range or long range nature of dependence in volatility still goes on (and may probably never be resolved),

agent-based models can provide motivation for choosing between alternative econometric specifications which are otherwise equally plausible in statistical terms, thus providing a useful complement to econometric analysis.

References

1. Andersen, T. and Bollerslev, T. (1997). Heterogeneous information arrivals and returns volatility dynamics. *Journal of finance*, **52**, 975–1005.
2. Arifovic, J. and Gencay, R. (2000). Statistical properties of genetic learning in a model of exchange rate. *Journal of Economic Dynamics and Control*, **24**, 981–1005.
3. Arthur, W., Holland, J., LeBaron, B., Palmer, J. and Tayler, P. (1997). Asset pricing under heterogeneous expectations in an artificial stock market, in: *The Economy as an Evolving Complex System*. Perseus Books, Reading MA, pp. 15–44.
4. Baillie, R.T., Bollerslev, T. and Mikkelsen, H.O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, **74**, 3–30.
5. Barndorff-Nielsen, O.E. and Shephard, 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.
6. Bayraktar, E., Horst, U. and Sircar, K.R. (2003). A limit theorem for financial markets with inert investors. *Working paper* Princeton University,
7. Beran, J. (1994). *Statistics for Long-Memory Processes*, Chapman and Hall, New York.
8. Bollerslev, T., Chou, R. and Kroner, K. (1992). ARCH modeling in finance. *Journal of Econometrics*, **52**, 5–59.
9. Breidt, F.J., Crato, N. and De Lima, P. (1998). The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics*, **83**, 325–348.
10. Cheridito, P. (2004). Gaussian moving averages, semimartingales and option pricing. *Stochastic Processes and Their Applications*, **109**, 47–68.
11. Chiarella, C., Gallegatti, M., Leombruni, R. and Palestrini, A. (2003). Asset price dynamics among heterogeneous interacting agents. *Computational Economics*, **22**, 213–223.
12. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, **1**, 1–14.
13. Cont, R., Bouchaud, J.-P. and Potters, M. (1997). Scaling in financial data: Stable laws and beyond, in: B. Dubrulle, F. Graner, and D. Sornette (eds.) *Scale Invariance and Beyond*, Springer, Berlin.
14. Cont, R., Ghoulmie, F. and Nadal, J.-P. (2004). Heterogeneity and feedback in an agent-based market model. *Working paper, Ecole Polytechnique*.
15. Cutler, D., Poterba, J. and Summers, L. (1989). What moves stock prices? *Journal of Portfolio Management*, **15**, 4–12.
16. Ding, Z., Granger, C.W.J. and Engle, R.F. (1993). A long memory property of stock market returns and a new model. *Journal of empirical finance*, **1**, 83–106.
17. Ding, Z. and Granger, C.W.J. (1996). Modeling volatility persistence of speculative returns: a new approach. *Journal of Econometrics*, **73**, 185–215.

18. Doukhan, P., Oppenheim, G. and Taqqu, M. S. (2003). *Theory and Applications of Long-Range Dependence*. Birkhäuser Boston Inc., Boston, MA.
19. Engle, R.F. (1995). *ARCH Models*. Oxford University Press, Oxford.
20. Giardina, I. and Bouchaud, J.-P. (2003). Bubbles, crashes and intermittency in agent based market models. *European Physical Journal*, **B. 31**, 421–437.
21. Giraitis, L., Leipus, R. and Surgailis, D. (2004). Recent advances in ARCH modelling. This volume.
22. Gouriéroux, C., Monfort, A. and Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, **8**, S85–S118.
23. Granger, C.W.J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics*, **14**, 227–238.
24. Granger, C.W.J. and Hyung, N. (2004). Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *Journal of Empirical Finance*, **11**, 399–421.
25. Guillaume, D., Dacorogna, M.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.
26. Heyde, C.C. (2002). On modes of long-range dependence. *Journal of Applied Probability*, **39**, 882–888.
27. Holland, J. (1992). *Adaptation in Natural and Artificial Systems*, Bradford Books.
28. Hommes, C.H., Gaunersdorfer, A. and Wagener, F.O. (2003). Bifurcation routes to volatility clustering under evolutionary learning. *Working paper, CenDEF*.
29. Kirman, A. (1993). Ants, rationality, and recruitment. *Quarterly Journal of Economics*, **108**, 137–156.
30. Kirman, A. and Teyssière, G. (2002). Microeconomic models for long memory in the volatility of financial time series. *Studies in Nonlinear Dynamics & Econometrics*, **5**, 281–302.
31. LeBaron, B. (2000). Agent-based computational finance : suggested readings and early research. *Journal of Economic Dynamics and Control*, **24**, 679–702.
32. LeBaron, B. (2001a). Evolution and time horizons in an agent-based stock market. *Macroeconomic Dynamics*, **5**, 225–254.
33. LeBaron, B. (2001b). Stochastic volatility as a simple generator of apparent financial power laws and long memory. *Quantitative Finance*, **1**, 621–631.
34. LeBaron, B., Arthur, B. and Palmer, R. (1999). Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, **23**, 1487–1516.
35. Liu, M. (2000). Modeling long memory in stock market volatility. *Journal of Econometrics*, **99**, 139–171.
36. Lobato, I. and Velasco, C. (2000). Long memory in stock market trading volume. *Journal of Business and Economic Statistics*, **18**, 410–427.
37. Lo, A. (1991). Long-term memory in stock market prices. *Econometrica*, **59**, 1279–1313.
38. Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tail of return distributions. *Journal of Economic Behavior and Organization*, **33**, 143–65.
39. Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: a micro simulation of interacting agents. *International Journal of Theoretical and Applied Finance*, **3**, 675–702.

40. Mandelbrot, B.B. (1963). The variation of certain speculative prices. *Journal of Business*, **36**, 392–417.
41. Mandelbrot, B.B. (1971). When can price be arbitrated efficiently? A limit to the validity of the random walk and martingale models. *Review of Economics and Statistics*, **53**, 225–236.
42. Mandelbrot, B.B. and Taqqu M.S. (1979). Robust R/S analysis of long-run serial correlation, in *Proceedings of the 42nd session of the International Statistical Institute*, Vol. 2 (Manila, 1979), **48**, 69–99.
43. Mandelbrot, B.B. and Van Ness, J. (1968). Fractional Brownian motion, fractional noises and applications. *SIAM Review*, **10**, 422–437.
44. Mézard, M., Parisi, G. and Virasoro, M. (1984). *Spin Glass Theory and Beyond*, World Scientific.
45. Mikosch, T. and Stărică, C. (2000). Limit theory for the sample autocorrelations and extremes of a GARCH (1, 1) process. *The Annals of Statistics*, **28**, 1427–1451.
46. Mikosch, T. and Stărică, C. (2003) Long-range dependence effects and ARCH modeling, in: Doukhan, P., Oppenheim, G. and Taqqu, M.S. (Eds.), *Long-Range Dependence: Theory and Applications*. Birkhauser, Boston, pp. 439–459.
47. Pagan, A. (1986). The econometrics of financial markets. *Journal of Empirical Finance*, **3**, 15–102.
48. Pourahmadi, M. (1988). Stationarity of the solution of $X_t = A_t X_t + 1$ and analysis of non-Gaussian dependent variables. *Journal of Time Series Analysis*, **9**, 225–239.
49. Resnick, S.I., Samorodnitsky, G. and Xue, F. (1999). How misleading can sample ACFs of stable MAs be? (Very!). *The Annals of Applied Probability*, **9**, 797–817.
50. Resnick, S.I. (1998). Why non-linearities can ruin the heavy-tailed modeler's day, in: Adler, P., Feldman, R. and Taqqu, M.S. (Eds.), *A Practical Guide to Heavy Tails: Statistical Techniques and Applications*, Birkhäuser Boston, Boston, MA, pp. 219–239.
51. Rogers, L.C.G. (1997). Arbitrage with fractional Brownian motion. *Mathematical Finance*, **7**, 95–105.
52. Taqqu, M.S. and Levy, J. (1986). Using renewal processes to generate long range dependence and high variability, in: Eberlein, E. and Taqqu, M.S. (Eds.), *Dependence in Probability and Statistics*, Birkhauser, Boston, pp. 73–89.
53. Teverovsky, V., Taqqu, M.S. and Willinger, W. (1999). A critical look at Lo's modified R/S statistic. *Journal of Statistical Planning and Inference*, **80**, 211–227.
54. Willinger, W., Taqqu, M.S. and Teverovsky, V. (1999). Long range dependence and stock returns. *Finance and Stochastics*, **3**, 1–13.
55. Zovko, I. and Farmer, J. D. (2002). The power of patience: a behavioral regularity in limit order placement. *Quantitative Finance*, **2**, 387–392.