

Bubbles, crashes and intermittency in agent based market models

I. Giardina^{1,2,a} and J.-P. Bouchaud^{3,4}

¹ Service de Physique Théorique, Centre d'études de Saclay, Orme des Merisiers, 91191 Gif-sur-Yvette Cedex, France

² Dipartimento di Fisica, Università di Roma La Sapienza and Center for Statistical Mechanics and Complexity, INFN Roma La Sapienza, Piazzale Aldo Moro 2, 00185 Roma, Italy

³ Service de Physique de l'État Condensé, Centre d'études de Saclay, Orme des Merisiers, 91191 Gif-sur-Yvette Cedex, France

⁴ Science & Finance, CFM, 109-111 rue Victor-Hugo, 92532 Levallois Cedex, France

Received 5 July 2002 / Received in final form 9 December 2002

Published online 14 February 2003 – © EDP Sciences, Società Italiana di Fisica, Springer-Verlag 2003

Abstract. We define and study a rather complex market model, inspired from the Santa Fe artificial market and the Minority Game. Agents have different strategies among which they can choose, according to their relative profitability, with the possibility of not participating to the market. The price is updated according to the excess demand, and the wealth of the agents is properly accounted for. Only two parameters play a significant role: one describes the impact of trading on the price, and the other describes the propensity of agents to be trend following or contrarian. We observe three different regimes, depending on the value of these two parameters: an oscillating phase with bubbles and crashes, an intermittent phase and a stable 'rational' market phase. The statistics of price changes in the intermittent phase resembles that of real price changes, with small linear correlations, fat tails and long range volatility clustering. We discuss how the time dependence of these two parameters spontaneously drives the system in the intermittent region. We analyze quantitatively the temporal correlation of activity in the intermittent phase, and show that the 'random time strategy shift' mechanism that we proposed earlier allows one to understand the observed long ranged correlations. Other mechanisms leading to long ranged correlations are also reviewed. We discuss several other issues, such as the formation of bubbles and crashes, the influence of transaction costs and the distribution of agents wealth.

PACS. 89.65.Gh Economics, business, and financial markets

1 Introduction

It is now well known that the statistics of price changes in financial markets exhibit interesting 'stylized facts', which are to some extent universal, *i.e.* independent of the type of market (stocks, currencies, interest rates, etc.) and of the epoch [1–4]. Price changes are in a good approximation uncorrelated beyond a time scale of the order of tens of minutes (on liquid markets). Their distribution is strongly non Gaussian: they can be characterized by Pareto (power-law) tails with an exponent in the range 3–5. Another striking feature is the intermittent nature of the fluctuations: localized outbursts of the volatility, *i.e.* the amplitude of the price fluctuations (averaged over a given time interval), can be identified. This fact, known as *volatility clustering* [2,3,5,6], can be analyzed more quantitatively: the temporal correlation function of the daily volatility σ_t can be fitted by an inverse power of the lag τ , with a rather small exponent in the range 0.1 – 0.3 [6–10].

This suggests that there is no characteristic time scale for volatility fluctuations: outbursts of market activity can persist for short times (a few hours), but also for much longer times, months or even years. The slow decay of the volatility correlation function leads to a multifractal-like behaviour of price changes [10–14], and has important consequences for option pricing. Other stylized facts have been reported, such as the leverage effect that leads to skewed distribution of price changes [15], or the apparent increase of inter-stock correlations in volatile periods.

It is now very clear to many that these features are very difficult to explain within the traditional framework of 'rational expectations', where all agents share the same information, have an infinite computation power and act in a perfectly rational way (see *e.g.* the clear discussion in the introduction of references [16,17] and in [18]). Another route, much less formalized and still very much in an exploratory stage, is followed by an increasing number of academics. The aim is to assume as little as possible about agents preferences and abilities, and to

^a e-mail: irene.giardina@roma1.infn.it

explore *generic* classes of models, with the hope of finding some *plausible* mechanisms that reproduce at least part of the stylized facts recalled above. In this endeavor, one should not be constrained by preexisting prejudices or established frameworks. The ‘grand unification’ of different mechanisms which would lead finally to a logically consistent and simultaneous understanding of all the empirical facts is deferred to later times. Similarly, it is premature to ask for rigorous proofs, but leave space for hand waving arguments and numerical simulations.

In this paper, we report [19] the results of an artificial market that bears some similarities with many previous attempts [20–38]. Although the detailed behaviour of our artificial market depends on the value of the different parameters entering the model, only a few market typologies are identified, and some qualitative features (such as the long range volatility correlations) are robust to parameter changes, at least in some regions of parameter space. We explain in particular that the general mechanism proposed in [39] is indeed responsible for volatility clustering in our model. We also discuss, in this context, other agent based models proposed so far in which this phenomenon has been observed. We identify several other interesting features (appearance of bubbles and crashes, influence of the price fixing procedure and of wealth constraints, of transaction costs, etc.) that may be of some relevance to real markets.

In our model, we have not included herding, or imitation effects. Each agent acts independently of other agents. The correlations between their actions is entirely mediated by the price history itself. Direct herding might also be important to account for the phenomenology of real markets (see [29, 40–42] and references therein). We leave to a future study the extension of the present model to account for these herding effects.

The aftermath of our study (and of all other similar studies) is the following paradox: in order to get a ‘good looking’ price chart, one has to tune quite a bit the important parameters of the model. What are the mechanisms tuning the parameters in real markets to make them look all alike? There must be some generic self-organization mechanisms responsible for this selection. We discuss in Section 6.1, in view of our results, what could be the ‘evolutionary’ driving forces relevant for this fundamental issue.

2 Set up of the model

2.1 Basic ingredients

In line with the original idea of the Santa Fe artificial market [21, 22], which was later simplified and popularized as the Minority Game (MG) [24, 43, 30, 44, 45], in our model agents do not follow a rational expectation paradigm but rather act inductively, adapting their behaviour to their past experience. As in the MG each agent has a certain fixed number of strategies, each of which converts some *information* into a *decision*. We will assume a world where

there are only two tradable assets: a stock, with fluctuating price, and a bond, yielding a certain (known) risk free rate ρ . The information on which the agents decide their action will be the past history of the price itself. The decision is to buy stocks (converting bonds in cash), to sell stocks, or to be inactive (*i.e.* to hold bonds). Each strategy is given a score, which is updated according to its performance. The strategy played at time t by a given agent is the one, among those available to him, which would have best performed in a recent past. We take proper account of the wealth balance of each agent, and proper market clearing (*i.e.* matching supply and demand) is enforced.

2.2 Notations and definitions

Each agent i , $i \in \{1, \dots, N\}$ has $S - 1$ active strategies plus an inactive one. He owns, at time t , a number $\phi_i(t)$ of stocks and $B_i(t)$ of bonds. The price of the stock is $X(t)$, and therefore the total wealth of agent i is $B_i(t) + \phi_i(t)X(t)$. The dynamics of the model, between t and $t+1$, is defined by the following set of rules:

– Information:

We assume that all agents rely on the *same* information \mathcal{I}_t , given by the m last steps of the past history of the return time series (m is for ‘memory’). We choose the information to be qualitative and to only depend on the *sign* of the previous price changes:

$$\begin{aligned} \mathcal{I}_t &= \{\chi(t-m), \dots, \chi(t-1)\} \\ \chi(t) &= \text{sign} \left[\log \left(\frac{X(t)}{X(t-1)} \right) - \rho \right]. \end{aligned} \quad (1)$$

In this sense, our traders are ‘chartists’ on short time scales, and take their decision based on the past pattern of price changes. (We will add below the possibility for the agents not to follow their systematic strategies, and act as ‘fundamental’ traders if the price is too far from a reference value, or even act at random.)

A comment on the value of the time scale ‘1’ is in order here. Clearly, different agents observe the price time series on different time scales, from several minutes for intra day traders to months for long term pension funds. Here, we assume for simplicity that all agents coarse-grain the price time series using the same clock, and consider as meaningful price variations on – say – a day or a week. All the parameters below were chosen such as one time step roughly corresponds to a week of trading. One interesting outcome of our model is that even if all agents have the same intrinsic clock, a broad range of time scales is spontaneously generated.

– Strategies:

Each agent i is endowed with a certain number S of fixed strategies, that convert information \mathcal{I}_t into decision $\epsilon_i(\mathcal{I}_t) = \pm 1, 0$ (buy, sell, inactive). For example, a ‘trend following’ strategy could be to choose $\epsilon_i = 1$ as soon as there is a majority of +’s in the signal \mathcal{I}_t . Another example is the inactive strategy, for which $\epsilon_i \equiv 0, \forall \mathcal{I}$.

Each agent is given the possibility to remain inactive, *i.e.* has an inactive strategy. The $S - 1$ other strategies are chosen at random in the space of all strategies (there are 2^{2^m} of them), in order to model heterogeneity in the agents capabilities. One can however give a bias to this random choice, and favor ‘trend following’ or ‘contrarian’ strategies. This is done by defining the ‘magnetization’ $M \in [-1, 1]$ of the string $\chi(1), \dots, \chi(m)$, defined as:

$$M = \frac{1}{m} \sum_{j=1}^m \chi(j), \quad (2)$$

and choose the corresponding decision to be $\epsilon = +1$ with probability $(1 + PM)/2$ and $\epsilon = -1$ with probability $(1 - PM)/2$. The parameter $P \in [-1, 1]$ can be called the ‘polarization’ of the strategies. If $P = 1$, trend following strategies are favored, whereas $P = -1$ corresponds to contrarian strategies. The choice $P = 0$ means no bias in the strategy space.

– Decision and buy orders:

Knowing the strategy used by agent i and the information \mathcal{I}_t allows one to compute the decision $\epsilon_i(t)$ of each agent. Depending on the value of $\epsilon_i(t)$, the agent buys/sells a quantity $q_i(t)$ proportional to his current belongings. More precisely, we set:

$$\begin{aligned} q_i(t) &= g \frac{B_i(t)}{X(t)} & \text{for } \epsilon_i(t) = +1 \\ q_i(t) &= -g \phi_i(t) & \text{for } \epsilon_i(t) = -1 \\ q_i(t) &= 0 & \text{for } \epsilon_i(t) = 0. \end{aligned} \quad (3)$$

This means that we consider ‘prudent’ investors who change their positions progressively: only a fraction g of the cash is invested in stock between t and $t + 1$ if the signal is to buy, and the same fraction g of stock is sold if the signal is to sell. Typical values used below are $g \sim 1\%$.

The normalized total order imbalance, which will be used to determine the change of price, is denoted $Q(t)$:

$$Q(t) = \frac{1}{\Phi} \sum_{i=1}^N q_i(t) = Q^+(t) - Q^-(t), \quad (4)$$

where Φ is the total number of outstanding shares (that we assume to be constant), Q^+ is the fractional volume of buy orders and Q^- the fractional volume of sell orders.

Agents sometimes choose to abandon their ‘chartist’ strategies when the price reaches levels that they feel unreasonable: when the price is too high, they are likely to sell, and *vice versa*. More precisely, we construct a long term average of past returns as:

$$\begin{aligned} \bar{r}(t) &= \frac{1 - \alpha}{\alpha} \sum_{t' < t} \alpha^{t-t'} r(t') \\ r(t') &= \log \left(\frac{X(t'+1)}{X(t')} \right), \end{aligned} \quad (5)$$

where $\alpha < 1$ defines the time scale $T_0 = 1/\log(1/\alpha)$ over which the averaging is done, and $r(t)$ is the instantaneous stock return. When \bar{r} is larger than a certain reference return ρ_0 , related to economy fundamentals, the stock can be deemed as overvalued and ‘fundamentalists’ will sell ($\epsilon_i = -1$). Conversely, if $\bar{r} < \rho_0$, the stock is possibly undervalued, and $\epsilon_i = +1$. We model the occurrence of fundamental trading as stochastic, by assigning a certain probability p_f for every agent to follow a fundamental strategy rather than a technical (chartist) strategy. We want p_f to increase with $|\bar{r} - \rho_0|$, and have chosen the following simple relation:

$$p_f = \min \left(1, f \frac{|\bar{r} - \rho_0|}{\rho_0} \right), \quad (6)$$

where f is a certain parameter describing the confidence of agents in fundamental information. Since p_f increases when the price goes up too fast, fundamentalists have a stabilizing role and give to the price a mean reverting component. In the following, we will assume that on the long run, the overall economy growth ρ_0 and the interest rate ρ are equal, and impose $\rho_0 = \rho$, although in practice the two fluctuate with respect to each other.

In the above rule, we have again assumed that all agents use the same time scale to determine the past average trend. This is probably very far from reality, where one expects that this time scale could be very different for different agents.

Finally, it can be useful to consider the influence of ‘irrational’ traders, who take their decisions on the basis of random coin tossing only. We define p_i as the probability for an agent to take a random decision. In this case, the probability to be a fundamentalist is $(1 - p_i)p_f$.

– Price formation and market clearing mechanism:

Once the aggregate order imbalance $Q(t)$ is known, we update the price following a simple linear rule [20, 21, 26, 27, 41]:

$$r(t) = \frac{Q(t)}{\lambda}, \quad (7)$$

where

$$r(t) = \log \left(\frac{X(t+1)}{X(t)} \right) \simeq \frac{X(t+1)}{X(t)} - 1 \quad (8)$$

is the logarithmic price return and λ is a measure of the ‘stiffness’ of the market. There has been recent empirical studies of this relation, which was shown to hold for individual stocks for small enough Q , on a sufficiently large time interval [47]. For larger order imbalance, the price response appears to bend downward, a possible consequence of the structure of the order books. We have included this effect, with no noticeable effect on the qualitative results presented below.

From a more microscopic point of view, *i.e.* on time scales smaller than the time unit that we have chosen, agents place orders of different types in the market: market orders and limit orders. Market orders allows the order to be executed with certainty, but at

the current market price. Limit orders ensures a maximum price for buy orders (and a minimum price for sell orders) but can be unexecuted, or only partially executed, depending on the history of the price. Therefore, in general, the order put down by an agent will be only partially filled. We assume that the fraction of unfulfilled orders is the same for all agents trying to buy (if the number of sold shares is insufficient), or for all agents trying to sell in the opposite case.

Market clearing is then ensured by the following rule: the global amount of sell orders is $Q^-(t)$, and the total number of shares that can be bought at price $X(t+1)$ is:

$$\tilde{Q}^+(t) = Q^+(t) \frac{X(t)}{X(t+1)}. \quad (9)$$

The fraction of filled buy orders φ_+ (resp. filled sell orders φ_-) is therefore:

$$\varphi_+ = \min\left(1, \frac{Q^-}{\tilde{Q}^+}\right) \quad \varphi_- = \min\left(1, \frac{\tilde{Q}^+}{Q^-}\right). \quad (10)$$

From these quantities, one determines the actual number of shares $\delta\phi_i$ bought or sold by agent i :

$$\begin{aligned} \delta\phi_i(t) &= g\varphi_+ \frac{B_i(t)}{X(t+1)} & \text{for } \epsilon_i(t) > 0 \\ \delta\phi_i(t) &= -g\varphi_- \phi_i(t) & \text{for } \epsilon_i(t) < 0. \end{aligned} \quad (11)$$

Note that φ_+ and φ_- are different: they are defined in a such a way that the total amount of sold shares is exactly the same as the total amount of bought shares at all times, as can be checked by inspection.

Of course, the above rules are to some extent arbitrary but are consistent and represent qualitatively, in a phenomenological way the results of a double-auction market in which the probability of being unexecuted is finite. Another rule, such that there are *no* unexecuted orders (and would correspond to a market with only market orders with a matching rule) would be to update the price in such a way that $\tilde{Q}^+(t) = Q^-(t)$. Interestingly, this leads to a very unrealistic dynamics where the price is not continuous but jumps up and down in a random way. This would indeed be the case in a market with no market makers and no limit orders: the lack of liquidity would lead to a totally unstable price process.

– Wealth dynamics:

We now have all the ingredients to update the number of stocks and bonds of agent i , *i.e.*:

$$\begin{aligned} \phi_i(t+1) &= \phi_i(t) + \delta\phi_i(t) \\ B_i(t+1) &= B_i(t)(1+\rho) - \delta\phi_i(t)X(t+1), \end{aligned} \quad (12)$$

where the last line adds the interest gained on the bonds between t and $t+1$ to the cash need to finance new stocks, or gained through stock selling. Note that there is an *injection* of wealth due to the positive interest rate ρ . We will discuss this further in the following.

– Update of the scores:

Each agent assigns scores to his strategies to measure their performance and uses at time t the best strategy, *i.e.* the one with highest scores. We need now to specify how the scores of the different strategies are updated. The score of the α th strategy of agent i at time t is denoted $\mathcal{S}_i^\alpha(t)$, whereas the decision associated to this strategy when the available information at time t is \mathcal{I}_t is $\epsilon_i^\alpha(\mathcal{I}_t)$. When the agent i decides at time t to trade, the price at which the trade takes place is $X(t+1)$. Therefore, the *virtual* profit he makes due to this trade is only known at time $t+2$ and is $\epsilon_i^\alpha(\mathcal{I}_t)[X(t+2) - X(t+1)]$. We choose to update the score of the active strategies proportionally to the relative profit, corrected by the interest rate¹:

$$\begin{aligned} \mathcal{S}_i^\alpha(t+1) &= (1-\beta)\mathcal{S}_i^\alpha(t) + \beta\epsilon_i^\alpha(\mathcal{I}_{t-1})[r(t) - \rho], \\ \alpha &= 1, \dots, S-1, \end{aligned} \quad (13)$$

whereas the score of the inactive strategy is identically zero. The parameter $\beta \leq 1$ defines a memory time: the performance of the strategies is only computed using the recent part of the history. More precisely, the performance is evaluated over a period $T \sim 1/|\log(1-\beta)|$ and is longer the smaller the value of β . The shortest memory corresponds to $\beta = 1$, when scores are evaluated on the basis of the last performance of the strategy performance².

Note that the update of the scores is not weighted by the actual transaction volume: good decisions are valued independently of the current wealth of the agent. Therefore, the score of the strategy is not proportional to the actual profit and loss curve.

One should keep in mind that for any $\beta \neq 0$ only the best strategy $\alpha^*(t)$ is played by the agent at time t . Nevertheless, he updates the scores of all strategies *as if* they had been played. In other words, market impact is neglected here, since the very fact of using a given strategy influences the price itself. The history of the price would have been different if a different strategy had been played. We do not take into account the market impact for two reasons: first, the update of the score is delayed as compared to the action itself (see Eq. (13)) – therefore, the main source of systematic bias discussed in the context of the Minority Game in, *e.g.* [46], is removed. Second, market impact is in practice very hard to estimate for traders themselves (although some recent studies start addressing this issue [47–49]), and strategies are often backtested under the assumption that the market impact is small.

A related point is that of the virtual profit computed above. Taking profit means closing one's position, at a

¹ A similar rule for the update of scores was recently considered in [38].

² Following equation (13), for $\beta = 0$ all scores are equal and time independent. According to our algorithm, each agent chooses at random the strategy to use. The value $\beta = 0$ is pathological: for any finite (even if small) value of β there are finite differences among the scores of the various strategies, and consequently a best strategy can be identified.




Fig. 1. Typical price charts in the three regimes: oscillatory, intermittent, stable (efficient). The parameters are: $S = 3$, $m = 5$, $g = 0.005$, $f = 0.05$, $N = 1001$, $\beta = 10^{-2}$ and $1 - \alpha = 10^{-4}$. The top graph corresponds to $g/\lambda = 0.1$, the two bottom graphs to $g/\lambda = 0.6$.

price that is not known in advance. Again there will be some market impact and the actual price of the transaction is on average less than the current price. This effect is well known to active market participants and, as mentioned above, has been recently the subject of some studies. A way to model this is to add a transaction cost to the above update of the score, independent of whether one buys or sells. (This cost should also be taken into account in the above wealth balance.)

2.3 Summary of the parameters and main results

The model contains a rather large number of parameters: the interest rate ρ , the memory length used for technical trading m , the ‘polarization’ of strategies P , the fraction of invested wealth g , the time scale used to trigger fundamental trading α , the propensity of fundamental trading f and the fraction of irrational agents p_i , the stiffness of the market λ , and the memory time of agents when the score of the strategies are updated, β . As we shall see, the most important parameters are g/λ and the polarization P , which determine the qualitative behaviour of price changes. The other parameters influence the quantitative results, but do not in general modify the qualitative features, apart from special cases which will be explicitly discussed in the following sections. As the values of g/λ and P are varied, we observe the appearance of three qualitatively different regimes (see Fig. 1):

- An *Oscillatory Regime*, corresponding to ‘weak coupling’: $g/\lambda \lesssim 0.4$, and $P \geq 0$, where speculative

bubbles are formed, and finally collapse in sudden crashes induced by the fundamentalist behaviour. In this regime, markets are not efficient, and a large fraction of the orders is (on average) unfulfilled.

- An *Intermittent regime* ($g/\lambda \gtrsim 0.4$, $P \geq -|P_0|$) where the ‘stylized’ facts of liquid markets are well reproduced: the market is efficient (although some persistent or antipersistent correlations survive), the returns follow a power law distribution, and volatility clustering is present.
- A *Stable regime*, which arises if the polarization P is sufficiently negative (predominance of contrarian strategies). In this case, the fluctuations of the price are mild and mean reverting (see Fig. 1c), as one would expect in a ‘rational market’ where the trading price is always close to the fundamental price.

A (rather schematic) phase diagram of the model in the plane $(g/\lambda, P)$ for a fixed value of all the other parameters is shown in Figure 2. This qualitative phase diagram is the central result of our study.

3 Kinematics of the model: fully random strategies

Before embarking to analyze the influence of strategies, it is important to calibrate the bare version of our model where agents take purely random, uncorrelated decisions at each instant of time. In this case, the price fluctuations will only reflect the wealth constraints. We show in




Fig. 2. Phase diagram of the model. The region $g/\lambda > 0.4$, $P > -|P_0|$ corresponds to an intermittent regime, with small linear correlations but strong volatility fluctuations. The dashed lines correspond to crossover regions, where a mixed behaviour is observed.




Fig. 3. Behaviour of the price as a function of time for purely random strategies. Inset: Variogram of the price fluctuations, and Ornstein-Uhlenbeck fit.

Figure 3 the price chart for some values of the parameters. The log-price performs a mean reverting random walk around the fundamental price $X_f(t) = \exp(\rho t)$. In the inset we show the log-price variogram, defined as:

$$\mathcal{V}(\tau) = \left\langle \left(\log \frac{X(t+\tau)}{X(t)} - \rho\tau \right)^2 \right\rangle, \quad (14)$$

together with an Ornstein-Uhlenbeck fit:

$$\mathcal{V}(\tau) = \mathcal{V}_\infty (1 - \exp(-\tau/\tau_0)), \quad (15)$$

that describes a mean reverting random walk with a reverting time τ_0 , and mean square excursion from the mean equal to \mathcal{V}_∞ . For small τ , the behaviour of $\mathcal{V}(\tau)$ is linear in τ , as for a free random walk, indicating that random trading leads, as expected, to unpredictable price changes. However, on larger time scales, the limitation

of wealth and of stocks (agents cannot borrow nor short sell stocks) prevents the price from wandering infinitely far from the fundamental price, and leads to a mean-reverting behaviour. This mechanism will also operate for more complicated trading rules and will be discussed again in Section 5.

It is simple to understand how these quantities depend on the parameters g, λ . From the price fixing mechanism, one can write a Langevin equation for the price that reads:

$$r(t) = \frac{d \log X}{dt} = \frac{g}{2\lambda\phi} \left(\frac{B(t)}{X} - \phi + \xi(t) \right), \quad (16)$$

where $B(t)$ is the average of an agent wealth in bonds, ϕ the average number of stocks per agent, and ξ is a random binomial variable measuring the small offset from a perfect 50–50 division between buyers and sellers. The variance of $\xi(t)$ is therefore of order N^{-1} . Equation (16) indeed describes a mean reverting process around $B(t)$: when the price is too high, the demand goes down due to budget constraints and the price goes down, and *vice versa*. The long term increase of price is only due, in our model, to the continuous injection of cash through the interest rate.

Equation (16) defines a mean reverting process, that allows one to obtain in particular:

$$\tau_0 \propto \frac{\lambda\phi}{g} \quad \mathcal{V}_\infty \propto \frac{g}{N\lambda\phi}, \quad (17)$$

in agreement with numerical results. The short time volatility of the market, for $\tau \ll \tau_0$, is given by $\sigma^2 = \mathcal{V}_\infty/\tau_0 \propto g^2/N\lambda^2\phi^2$, and is small for ‘stiffer’ markets, or if the fraction of invested wealth is smaller, as expected. Also, the volatility decreases when the number of agents increases. For a market with 10^4 participants such that $g = 10\%$, this formula gives a reasonable volatility of 1% per week, when the market stiffness is $\lambda \sim 0.1$. However, in this case, the return time is also of the order of a week and the total variability of prices is 1%, both being far too small compared to reality.

Let us finally note that from the simulations, the market liquidity, measured as the fraction c of fulfilled orders, improves when the market stiffness λ decreases. For example, we find $c = 0.9$ for $g = 1\%$, $\lambda = 1$, and $c = 0.98$ for $g = 1\%$, $\lambda = 0.1$.

4 The oscillatory regime

This regime is characterized by the presence of regular bubbles followed by rapid ‘crashes’. The period of the bubbles is a function of the model parameters. In Figure 4 we show the dependence of the period on g/λ , for some fixed values of the other parameters. The period decreases as g/λ increases, and vanishes when the market enters the intermittent regime. On the same plot, we have also shown the fraction c of fulfilled orders, which is very low in the oscillatory phase and increases with g/λ .




Fig. 4. Period T of the oscillations (open circles, left scale), and fraction c of fulfilled orders (black squares, right scale), as a function of g/λ . All other parameters are as in Figure 1.

4.1 A qualitative discussion

We want to understand the mechanism underlying the creation and persistence of bubbles, and their final collapse.

To answer this question, one has to look back at equation (7) that describes the price dynamics. Since the strategies are randomly distributed among agents, we can expect that at the beginning of the game one half of the agents is willing to buy and the other half to sell. The price increment is therefore $r = \delta X/X \sim g\lambda^{-1}N/(2\Phi)[B/X - \Phi]$. Since in general $B/X \neq \Phi$, the price rises or decreases. Suppose that the initial conditions are such that it increases. If g/λ is small, both the wealth and number of shares of the agents change quite slowly, and therefore $B/X - \Phi$ will keep a constant sign for a while. This will generate a history of price of the form $\dots, 1, 1, 1, 1, 1, \dots$ (see Eq. (1)). This is the initial stage of the bubble. Now, the population of agents fall into two categories: those who have at least one strategy such that $\dots, 1, 1, 1, 1, 1$ sends a buy signal, and those who do not have such a strategy available. The buy-strategies keep being rewarded, while the sell ones keep losing points: agents who can buy will continue buying, while agents who cannot will soon become inactive since all active strategies have a negative score. In this way the bubble is self-sustained and the price keeps increasing. This ‘majority’ mechanism was discussed in the context of the Minority Game in [36,38]. However, the buying power of the buyers decreases, because the price increases and the available cash decreases. Hence, the relative return of the stock over the risk-free rate diminishes, and the score of the buying strategies only become marginally positive (because the initial large gains are forgotten, only the recent past is included in the calculation of the scores). Since the average past return exceeds the reference rate ρ_0 , a certain fraction of agents become fundamentalist and act contrarily to the main trend, *i.e.* sell. As soon as their action is such that the price drops, the history will change to $\dots, 1, 1, 1, 1, -1$.

At this point, one half of the active agents still receive a buy signal, but the other half receive a sell signal (since strategies are random, and for now unbiased – $P = 0$). Now, these selling agents have a lot of stocks, since they have been buying for a long time, whereas the still buying agents have a poor buying power. This obviously results in a series of -1 , and an ‘anti-bubble’ is created. This anti-bubble has a large initial negative slope because the system starts in a highly unbalanced state. There is no symmetry between a bubble and an antibubble because of the presence of a non zero interest rate, which is a source of cash and favors, on average, positive trends. Once again, it is the presence of fundamentalist agents which at some point triggers the end of the descending trend, and re-establish a bubble. The precise time at which the bubble collapses is random, and the way the collapse is triggered is very similar to a nucleation process (see next section).

One interesting quantity is the fraction of fulfilled orders. In the oscillating regime, one understands from the above arguments that the unbalance between buy orders and sell orders is in general very large, leading to great amount of unfulfilled orders. The fraction c of fulfilled orders increases with g/λ : see Figure 4. As will be discussed below, this is a driving force to escape from this oscillating regime, which obviously does not look at all like real markets.

The effect of the polarization P , which can be appreciated in the phase diagram of Figure 2, is quite easy to understand in the light of this microscopic analysis: a positive value of P increases the average number of trend followers and therefore is almost irrelevant for the bubble dynamics which is based on an endogenous trend following behaviour. On the other hand $P < 0$ acts contrarily to the bubble creation forcing a percentage of agents to act against the trend. Interestingly enough, as it can be seen from the phase diagram, even a small $P < 0$ is able to prevent the appearance of a bubble. As a function of $|P|$, the transition between an oscillating behaviour and a stable behaviour is first-order like, in the sense that the period of the oscillation is still finite when the transition occurs.

The above mechanism also allows us to understand the role of the parameters f or the fraction of irrational agents p_i . The numerical simulations indicate that increasing the value of these parameters may at first *stabilize* the oscillating phase, while one would intuitively expect an opposite behaviour since fundamentalists/irrational players reduce the relative number of trend followers. This effect can be explained in the following way. We have seen that, as the price increases during the bubble, the available cash of active players and therefore their buying power decreases bringing the bubble to saturation. However, each active agent has a certain probability proportional to f to become a fundamentalist and sell, increasing his wealth in cash. The buying pressure can then stay higher for a longer period of time, determining more stable bubbles. Of course, as f increases substantially, its other effect of nucleating opposite trends becomes the most important one, and the system behaves much as if the polarization parameter P was strongly negative, and thus prevents the

appearance of bubbles. Finally, the presence of irrational agents who buy or sell randomly, and therefore statistically drop out of the order imbalance, can be seen as reducing the effective value of g , and thus – somewhat surprisingly – stabilizing the oscillating regime. On the other hand, in this regime the properties of the system behaviour do not depend significantly on the memory β . Indeed, during the bubble even a very short memory is enough to select the buying strategy as the best one because buying strategies win at each step of time. A similar argument holds for selling strategies during the crash.

We now turn to a mathematical transcription of the above discussion, that allows us to characterize more accurately the *shape* of the bubbles in our model.

Note finally that, when each agent has a small number of strategies, the oscillating regime tends to die out for very long times, since the amplitude of the bubbles progressively decreases. This is due to the fact that agents that have a buying strategy during the bubble tend to underperform on the long run. Their wealth is, at long times, insufficient to sustain the bubble. In order to determine the oscillation period more precisely, we have artificially given to each agent both its initial strategy and its perfect mirror image. This prevents the appearance of two groups with systematically different wealths, without changing the basic mechanism leading to bubbles and crashes.

4.2 A mean-field description

In this section we try to describe the bubble-crash dynamics by means of some mean-field equations. As we shall see, this will enable us to understand the precise mechanism of the bubble saturation and the consequent crash occurrence.

We start by defining the relevant mean-field variables: the average amount B of bonds and the average stock amount ϕ hold by active agents. It is also useful to define the number of buyers/sellers N_b, N_s . Note that since only part of the agents are active during the bubble-crash dynamics ϕ is different from Φ/N , the average stock amount per agent. Besides, since active players may act as trend followers but also as fundamentalists, ϕ may vary in time.

Our aim is to write some evolution equations for these average quantities and for the price. In the following we shall consider the continuous time limit of equations (7), (12) and (5). We have:

$$r(t) = \frac{d \log X}{dt} = \frac{1}{\lambda} Q(t) \quad (18)$$

$$\bar{r}(t) = \frac{1}{T_0} \int_0^t d\tau \exp\left(-\frac{t-\tau}{T_0}\right) r(\tau), \quad (19)$$

where $T_0 = 1/\ln(1/\alpha)$ (see Eq. (5)). On the other hand the precise expression of the global action $Q(t)$ as well as the wealths dynamics is not the same during the whole bubble-crash cycle and to write down appropriate equations it is necessary to distinguish different regimes where

the agents behaviour can be considered as homogeneous in time. In general we can assume that during the bubble all the active players are buyers, except for a small percentage of fundamentalists who start selling when the price increases too much. On the contrary, during the crash all active agents become sellers while the fundamentalists will eventually act as buyers. In the following, we shall consider the case where the parameters g/λ and f are very small.

4.2.1 Beginning of the bubble: $\bar{r}(t) < \rho$

At the beginning of the bubble the price starts increasing but is still lower than the fundamental reference value $\exp(\rho t)$. This implies that all active agents, including the fundamentalists, are buyers. The total number of buyers is therefore given by all the agents who have at least one buying strategy, plus a stochastic contribution given by the fundamentalists:

$$N_b = N \left(1 - \frac{1}{2^s} \right) + \frac{N}{2^s} p_f, \quad (20)$$

where $s = S - 1$. On the other hand, the active sellers rapidly disappear as the score of the corresponding strategy deteriorates. The price rises without transaction, but since g/λ is small, it does so rather slowly. since all players who have not a buying strategy prefer to remain inactive. In this context, it is evident that no transaction can be realized at all since there is a complete unbalance between offer and demand, and the parameter φ_+ , which determines the fraction of fulfilled orders, is therefore identically zero. The evolution equations are in this case very simple (see Eqs. (7, 12)):

$$r(t) = \frac{g}{\lambda} N_b \frac{B}{\Phi X} \quad (21)$$

and:

$$\frac{dB}{dt} = \rho B \quad \frac{d\phi}{dt} = 0. \quad (22)$$

The initial rate of growth of the bubble r_0 is therefore given by:

$$r_0 \approx \frac{g}{\lambda} \frac{B_0 N}{X_0 \Phi} \left(1 - \frac{1}{2^s} \right), \quad (23)$$

where $B_0 = B(t = t_0)$ and $X_0 = X(t = t_0)$, and t_0 is the start date of the bubble. This gives a good approximation, as can be seen from Figure 5. Since the value of X_0 just after the crash is small, the buying power is high and r_0 is much larger than ρ . The value of $\bar{r}(t)$ therefore steadily increases.

4.2.2 Saturation of the bubble: $\bar{r}(t) > \rho$

In this regime the fundamentalist players start acting contrarily to the main trend and take their profit, thereby reestablishing some trading activity. We now have:

$$N_b = N \left(1 - \frac{1}{2^s} \right) (1 - p_f) \quad N_s = N p_f \quad (24)$$




Fig. 5. Bubble-crash dynamics in the oscillatory regime for $g/\lambda = 0.005$, and $\rho = 0.001$, $f = 0.005$, $\alpha = 0.9999$. The straight line corresponds to numerical data, the dashed lines to the mean-field approximations (see text), and the dotted line to the best exponential fit in the bubble saturation regime.

and the return and wealths evolution equations now read

$$r(t) = \frac{g}{\lambda \Phi} \left(N_b \frac{B}{X} - \frac{N_s}{N} \Phi \right) \quad (25)$$

$$\frac{dB}{dt} = \rho B - (1 - p_f) g \varphi_+ B + p_f g \phi X \quad (26)$$

$$\frac{d\phi}{dt} = (1 - p_f) g \varphi_+ \frac{B}{X} - p_f g \phi \quad (27)$$

where we have assumed that $\varphi_+ = p_f \Phi X / N_b B \ll 1$, which is consistent when $f \ll 1$. In this limit, the equations simplify significantly, since all terms containing p_f can be neglected. Introducing the buying power $\mathcal{Y} = B/X$, one has:

$$r \approx \frac{gN}{\lambda \Phi} \left(1 - \frac{1}{2^s} \right) \mathcal{Y} \quad \frac{d\mathcal{Y}}{dt} \approx (\rho - r) \mathcal{Y}, \quad (28)$$

from which one extracts:

$$\frac{dr}{dt} = (\rho - r)r \longrightarrow r(t) = \frac{\rho}{1 - A \exp(-\rho t)}, \quad (29)$$

with $A = (r_0 - \rho)/r_0 > 0$. This shows that $r(t)$ tends for long times to the overall growth rate of the economy. This is expected since a faster growth cannot be sustained because of budget constraints: the buying power would then tend to zero. We also find that $r(t)$ tends to ρ exponentially, with a decay rate γ equal to ρ itself. A better approximation for γ can be obtained by retaining the small terms dropped in the above analysis. In Figure 5 we show the curve obtained from equation (29) with the value of γ estimated analytically together with the best exponential fit of the data, and find very good agreement.

In this limit where $\rho T_0 \gg 1$, one can show that $\bar{r}(t) - \rho$ hardly varies during the bubble period, and is approximately equal to $A\rho$. We show in Figure 6 the time evolution of various quantities during this bubble.




Fig. 6. Upper graph: Behaviour of $r(t)$ (note the fluctuations) and $\bar{r}(t)$ as a function of t in the bubble-saturation regime (respectively, straight and dashed lines). Lower graph: For the same time-interval, the buying power B/X (straight line) and ϕ (dashed line). Same values of parameters as in Figure 5.

The knowledge of γ also allows us to estimate the time of occurrence of the crash. We know that the bubble is sustained because buying strategies keep being rewarded when the price keeps increasing. Nevertheless, we have seen that during the bubble saturation the average return approaches the reference value ρ , thus also the performance of buying strategies (*i.e.* their scores relative to the risk free rate) will increase less as time goes on; a small drop of price might then be enough to both trigger some strong selling strategies and favor the inactive strategy, reducing the buying pressure. In the absence of fluctuations, the average role of the fundamentalists is to renormalize the values of A and γ in equation (29) above. The number of fundamentalists is however a stochastic Poisson process, and fluctuates around the average value Np_f with variance $\sigma_f = \sqrt{Np_f}$. One therefore expects that when $r(t) - \rho$ becomes of the order of $g\sigma_f/\lambda N$, one of such positive fluctuations will bring the instantaneous return below ρ and the price history will therefore change from 1, 1, 1, ..., 1 to 1, 1, 1, ..., -1, from which a sudden wave of sell orders triggers the crash. A reasonable estimate of the bubble end-time t^* is therefore given by:

$$\rho A \exp(-\gamma(t^* - t_0)) \approx \frac{g}{\lambda} \sqrt{\frac{p_f}{N}}, \quad (30)$$

or

$$t^* - t_0 \approx \frac{1}{2\gamma} \log \left(\frac{A\lambda^2 \rho^2 N \gamma T_0}{fg^2} \right) \quad (31)$$

which indeed has the correct order of magnitude and qualitative behaviour, increasing with decreasing g/λ (see

Fig. 4). Note also that as $f \rightarrow 0$, the bubbles have an infinite lifetime, as we have seen in our simulations. Some fundamentalist behaviour is needed to nucleate a crash.

4.2.3 Crash

At the beginning of the crash all the fundamentalists act as sellers since one still has $\bar{r}(t) > \rho$. Thus

$$N_b = 0 \quad N_s = N \left(1 - \frac{1}{2^s} \right) + Np_f \frac{1}{2^s} \quad (32)$$

and the initial slope of the crash is now given by:

$$r^* \approx -\frac{g}{\lambda} \frac{N}{\Phi} \left(1 - \frac{1}{2^s} \right). \quad (33)$$

When at some point $\bar{r}(t)$ becomes smaller than ρ the fundamentalists start acting again contrarily to the main trend driving back the price toward the fundamental value $X_f(t)$. We may expect a saturation effect as in the previous regime, which would enable an asymptotic expansion. However, as it can be seen from the figures, the crash period is much shorter than the bubble one and the crash stops much before any asymptotic regime set in. The reason for this asymmetry is the interest rate that ‘refills’ the buying power in the bubble regime, that has no counterpart in the crash regime. The crash ends in the same way as the bubble does, through a nucleation process when $r(t) - \rho$ becomes of the order of the fluctuations (which are in this case much larger).

5 The intermittent regime

5.1 Results from the simulation

Upon increasing the parameter g/λ , the oscillating regime disappears and gives rise to an interesting market behaviour, where different market ‘states’ coexist: bubbles and crashes, periods of very small activity, and periods of very large activity, intermixed with each another. A typical chart of the returns and of the volume as a function of time is plotted in Figure 7. From visual inspection, it is quite clear that the return time series exhibits volatility clustering. We now turn to a more quantitative analysis of the price series statistics. The data we analyze below corresponds to $N = 10\,000$ agents, with ‘unpolarized’ strategies $P = 0$, and for $g/\lambda = 0.75$, deep in the intermittent regime (see Fig. 2). Other parameters are identical to those in Figure 1. We note that the qualitative effects we report below have not been seen to depend on N , at least up to $N = 10\,000$ which is the largest size we have investigated. Strong size effects, reported in the Lux-Marchesi model [29] for instance, seem to be absent in our case. However, we expect that when N becomes comparable to the total number of strategies (*i.e.* 2^{2^m}), the phenomenology will change since many agents will share exactly the same strategies.




Fig. 7. Time series of the returns, truncated to $\pm 10\%$ (top panel), and of the fraction of active agents (bottom panel), for $g/\lambda = 0.75$.




Fig. 8. Variogram of the price as a function of the time lag, together with an Ornstein-Uhlenbeck (mean-reverting) fit.

First, we look at the price variogram, defined by equation (14). This is shown in Figure 8, together with an Ornstein-Uhlenbeck fit. The saturation time τ_0 is of the order of 100; taking the unit time in our model to be the week, this corresponds to one or two years, with a volatility of roughly 20% per year, which is quite reasonable. As explained in Section 3 the saturation is a direct consequence of the bounded wealth of agents. Is there a similar mean-reverting trend in real markets? The analysis of the variogram of the Dow-Jones index, for example, in the period 1950–2000, shows no convincing sign of saturation on the scale of the year, although a slight bend down wards

for longer time lag is visible, but the data become noisy. There has been reports in the literature of a systematic mean-reverting effects on the scale of 5–10 years, perhaps related to the mechanism discussed here. Remember that the fundamental price in our model has zero volatility. The short time volatility is the result of pure trading, which leads to a random walk like behaviour of the price; the saturation occurs because of insufficient resources to sustain a large difference between the fundamental price and the speculative price.

The two other quantities that we have systematically studied are the volume variogram and the absolute return variogram. The variogram of a time dependent quantity is defined as:

$$\mathcal{V}(\tau) = \left\langle (O(t + \tau) - O(t))^2 \right\rangle. \quad (34)$$

We have considered the cases where $O(t)$ is equal to the fraction of active agents (defining \mathcal{V}_v) and where $O(t) = |r(t)|$ (defining \mathcal{V}_σ). We show these two quantities in Figure 9 as a function of $\sqrt{\tau}$, together with a fit inspired from the theory explained in the next section:

$$\mathcal{V}(\tau)|_{SQRT} = \mathcal{V}_\infty \left(1 - \exp \left(-\sqrt{\frac{\tau}{\tau_0}} \right) \right). \quad (35)$$

Note that the short time behaviour of the above function is $\sqrt{\tau}$, in contrast with the regular (linear) behaviour of the Ornstein-Uhlenbeck form. We will explain the origin of this singularity below, and explain why equation (35) fits very well the volume variogram. Note also that $|r(t)|$ is only a noisy proxy for the true quantity of interest, namely the volatility. Since the ‘measurement’ noise is expected to be uncorrelated at different instants of time, we use a

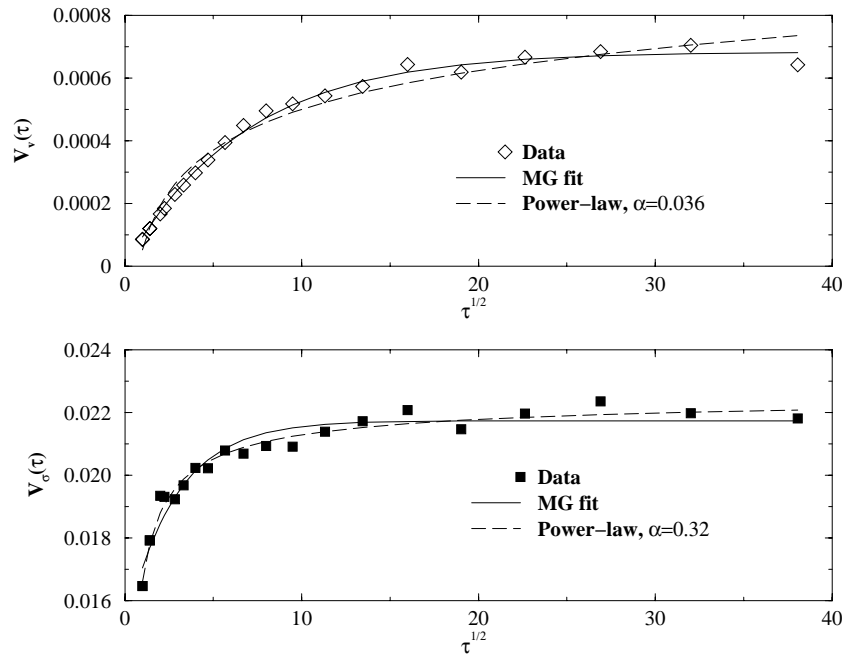


Fig. 9. Variogram of the volume (top panel) and of the absolute return (bottom panel) as a function of the time lag, together with the fits given by equations (35, 37). The parameters of the fit are given in the text.

slightly modified fit (which is also useful to deal with real market data):

$$\mathcal{V}(\tau)|_{SQRT} = \mathcal{V}_\infty \left(1 - \exp\left(-\sqrt{\frac{\tau}{\tau_0}}\right) \right) + \mathcal{V}_0 (1 - \delta_{\tau,0}). \quad (36)$$

We have also used a power-law fit

$$\mathcal{V}(\tau)|_{PWR} = \mathcal{V}_\infty \left(1 - \left(\frac{\tau}{\tau_0}\right)^{-\alpha} \right), \quad (37)$$

which has been advocated in many empirical studies of the volatility, with $\alpha \sim 0.1 - 0.3$ [8, 10].

As can be seen from Figure 9, the two fits are of comparable quality. The relevant parameters of the fit are, for \mathcal{V}_v , $\tau_0 \approx 46$ using equation (35) and $\tau_0 \approx 1$ and $\alpha \approx 0.0365$ using equation (37), and for \mathcal{V}_σ , one finds $\tau_0 \approx 7.4$ and $\mathcal{V}_0/\mathcal{V}_\infty \approx 2.2$ using equation (36) and $\tau_0 \approx 0.017$ and $\alpha \approx 0.325$ using equation (37). Note that the value of α found for $\mathcal{V}_v(\tau)$ is significantly smaller than that for $\mathcal{V}_\sigma(\tau)$, as also found for real market data. The reason for this will be explained in the next section.

We have also studied the distribution of returns. Not surprisingly, this distribution is found to be highly kurtic, which is expected since the volatility is fluctuating. The tail of the distribution can be fitted by a power-law, with an exponent $\mu \approx 3.5$, similar to the value reported in [50]: see Figure 10. Note that the negative tail is slightly fatter than the positive tail. We also show a Student-t distribution with a tail index $\mu = 4$, that fits quite well the whole distribution.

The role of a non zero polarization of strategies P is, for small enough P , to induce some correlations (or anticorrelations) in the returns. For $P < 0$, as was the case

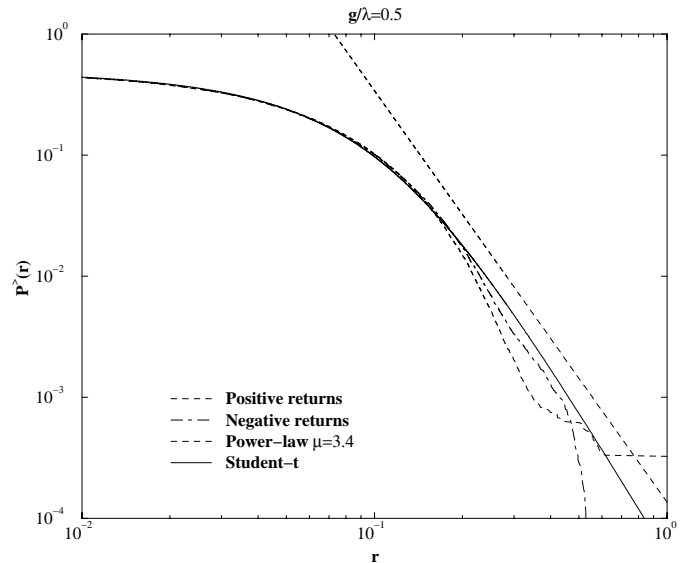


Fig. 10. Distribution of returns for $g/\lambda = 0.5$. A power-law tail with $\mu = 3.4$ and a Student-t distribution with a tail index $\mu = 4$ are shown for comparison.

for the oscillating regime, there is a first-order (discontinuous) transition toward a stable market, with small, strongly anticorrelated fluctuations that track the fundamental price. This transition occurs for rather small values of $|P| \approx 0.03$. For $P > 0$, on the other hand, the intermittent phase survives but the variogram of price fluctuations shows significant positive correlations. For sufficiently large P , the oscillating phase reappears in a continuous way.

The conclusion of this section is that volatility clustering appears for large values of g/λ . Qualitatively similar effects are seen for different choices of m (memory time for the strategies), β (memory time for the scores), and S (number of strategies per agent), provided g/λ is large enough to be in the intermittent phase. A crucial ingredient, however, is the existence of an inactive strategy, *i.e.*, the fact that the volume of activity is allowed to fluctuate. We have not been able to obtain long term volatility correlations of the type reported in Figure 9 when all strategies are active. This observation has motivated us to propose a simple mechanism for non trivial volume (and volatility) fluctuations [39,19], that we discuss now in the present context.

5.2 A simple mechanism for long-ranged volume correlations

5.2.1 Random time strategy shifts

In the above model, as in the Minority Game, scores are attributed by agents to their possible strategies, as a function of their past performance. In particular, the inactive strategy is adopted when the score of all active strategies are negative.

In the intermittent regime we have seen that the market is ‘quasi-efficient’: the autocorrelation of the price increments is close to zero. To a first approximation no strategy can on the long run be profitable. This implies that the strategy scores locally behave, as a function of time, as random walks. This very fact enables us to explain the fluctuations of volume. Let us consider for simplicity the case $S = 2$ (one active strategy and one inactive strategy per agent). Since the switch between two strategies occurs when their scores cross, the activity of an agent is determined by the survival time of the active strategy over the inactive one, that is by the return time of a random walk (the score of the active strategy) to zero. The interesting point is that these return times are well known to be power-law distributed. This leads immediately to the non trivial behaviour of the activity variogram shown in Figure 9. The same argument was used to explain the volume fluctuations in the Minority Game with an inactive strategy, in the efficient phase [39].

More formally, let us define the quantity $\theta_i(t)$ that is equal to 1 if agent i is active at time t , and 0 if inactive. The total activity is given by $v(t) = \sum_i \theta_i(t)$, and the activity variogram is given by

$$\begin{aligned} V_v(t, t') &= \left\langle [v(t) - v(t')]^2 \right\rangle \\ &= C_v(t, t) + C_v(t', t') - 2C_v(t, t'). \end{aligned} \quad (38)$$

where $C_v(t', t)$ denotes the volume correlation function.

One can consider two extreme cases which lead to the same result, up to a multiplicative constant: (a) agents follow completely different strategies and have independent activity patterns, *i.e.* $\langle \theta_i \theta_j \rangle \propto \delta_{i,j}$ or (b) agents follow very similar strategies, in which case $\theta_i = \theta_j$. In both

cases, $C_v(t, t')$ is proportional to $\langle \theta_i(t) \theta_i(t') \rangle$ and can be exactly expressed in terms of the distribution $P(s)$ of the survival times s of the active strategies [51,39]. For an unconfined random walk, the return time distribution $P(s)$ decays as $s^{-3/2}$ for large s . Note that this $s^{-3/2}$ behaviour is super-universal and only requires short range correlation in the score increments, not even a finite second moment [52,53]. However, in our model, the finite memory with which the scores are updated (*i.e.* the value of $\beta < 1$ in Eq. (13)) leads to a truncation of the $s^{-3/2}$ beyond a time $\tau_0 \simeq 1/\log(1 - \beta)$. Without this truncation, the volume $v(t)$ would never become a stationary process, *i.e.* $C_v(t, t')$ would still depend on both t and t' at long times, a phenomenon called ‘aging’ in the physics literature [54]. For finite τ_0 , one can show analytically that [39,19]:

$$V_v(t + \tau, t) = V_v(\tau) \propto \sqrt{\frac{\tau}{\tau_0}} \quad \tau \ll \tau_0. \quad (39)$$

The formula given by equation (35), was found to reproduce very well the crossover from this exact short time singular behaviour to the saturation regime.

The interesting conclusion is the following: the very fact that agents compare the performance of two strategies on a random signal leads to a multi-time scale behaviour of the volume fluctuations. This argument accounts very accurately for our numerical data both for the Minority Game and the present market model (see Fig. 9) [39,19], and also reproduces quite well the empirical variogram of activity in real markets [39].

5.2.2 Volume and volatility

Let us now discuss the relation between volume and volatility. In real markets, the two are known to be correlated; more precisely, it has recently been shown in [55] that the long run correlations in volatility come from the long range correlation in the volume of activity (see also [56]). In our artificial market, a scatter plot of the logarithm of the absolute return, $\log|r(t)|$ versus the volume of activity $v(t)$ shows nearly no correlations when the volume has small fluctuations around its average value \bar{v} , but is strongly correlated with $v(t)$ when $v(t)$ has large excursions above \bar{v} . This means that periods of high activity are also periods of large volatilities. One therefore expects that the structure of temporal correlations of the volume discussed above is also reflected in the volatility. If the relation between $|r(t)|$ and $v(t)$ was linear, or weakly non linear, one would in fact expect exactly the same shape for the variogram. The fact that this relation is highly non linear, *i.e.*, nearly no correlations for $v(t) \approx \bar{v}$, and a roughly exponential relation for $v(t) > \bar{v}$, adds an instantaneous noise contribution proportional to $1 - \delta_{\tau,0}$ to the variogram of absolute returns, and leads to a strong distortion of the shape of the relaxation. The fact that the effective power-law α , defined in equation (37), is larger for $|r(t)|$ than for $v(t)$ can be understood in details in the context of the multifractal random walk model of Bacry *et al.* [10].




Fig. 11. Variogram of the absolute returns as a function of the time lag in the Lux-Marchesi model with parameters as in [32]. We compare with the fit suggested by the random time strategy shift mechanism (Eq. (35)) and a power-law fit equation (37), as suggested by multifractal models [10]. Data: courtesy of Thomas Lux.

5.2.3 Other mechanisms for long-ranged correlations

Recently, many agent based models have been proposed to account for the stylized facts of financial markets, in particular volatility clustering and long range dependence [31, 32, 37, 57–60]. From the analysis of these models, one can distinguish three main mechanisms for this long-range dependence:

- Subordination of the strategies to performance. This is the mechanism explained above: as soon as each agent has different strategies with different levels of activities, and that the choice between these strategies is subordinated to their performance, one expects to see long range dependence of the type described above, whenever these strategies lead to identical long term performance. A similar mechanism is found in the models of [32, 57, 58], where agents switch between different trading styles (*e.g.* fundamentalists/chartists) as a function of their perceived performance and of herding effects. The basic prediction of this scenario is the short time square root singularity of the volume variogram. This prediction is very well obeyed in the Lux-Marchesi model, as shown in Figure 11.
- Subordination of the volatility to the price. In many models, the level of activity depends on the difference between the current price and a fundamental price. For example, in the model considered by Bornholdt [59], the volatility is a growing function of the absolute difference of these prices. In the model recently studied by Alfarano and Lux [60], on the contrary, the volatility is a decreasing function of this difference. In these

models, the price is mean reverting toward the fundamental price. Calling $y(t) = \log(X(t)/X_f(t))$, where X_f is the fundamental price, a schematic equation for $y(t)$ is:

$$\frac{dy}{dt} = -\kappa y + \sigma(y)\xi(t), \quad (40)$$

where ξ is a white noise and $\sigma(y)$ a certain function. The corresponding temporal correlation function of the volatility can, for some specific forms of $\sigma(y)$, be exactly computed, and generically leads to a non exponential decay that can mimic long term dependence. Empirically, one can indeed detect, on the Dow Jones index over a century, some correlation between the volatility and the difference between the current level and the average level of the index. However, we have checked using empirical data that that volatility clustering exists even when the level of market is neither particularly overvalued nor undervalued.

- Heterogeneity of the agents time scales. Another mechanism, close in spirit to the ‘HARCH’ model [61] or the cascade models proposed recently [9–13], comes from the different time horizons used by the agents to set up their strategies. For example, in the model of Raberto *et al.* [37], agents place orders at a distance from the current price proportional to a sliding average of the past volatility. The time scale used by the different agents is uniformly distributed between 10 day and 100 days. Correspondingly, this induces a non exponential decay of the volatility correlation function when $\tau \leq 100$.

All these three mechanisms are expected to play a role in real markets. More precise statistical tests will hopefully allow one to distinguish between them and estimate their relative contribution to the long-range dependence effects.

5.3 Crashes: the role of memory and the dynamical freezing of the choice mechanism

We have also observed the following two interesting effects in the intermittent phase:

- The frequency of bubbles/crashes decreases when the time horizon m of the strategies is increased.
- Just after a crash that follows a speculative bubble, the volatility is anomalously small.

First, how do bubbles form in this intermittent regime? The price fluctuations are almost without correlations. In this context it happens with probability $1/2^{m+1}$ that at a certain time step the m -bit past price history is $(1, 1, \dots, 1)$ and simultaneously the buying power of buyers is larger than that of sellers. In this case, at the next time step the price again increases and the past history is mapped onto itself. Also, since g/λ is large, the price change is substantial and the score of the buying strategies increases fast, thereby creating a bubble that will terminate in the same way as we described in Section 4. Thus, we expect an average frequency of $1/2^{m+1}$ for the bubble/crash occurrences, which one can easily verify in our model.

The second question concerns the period of very low volatility which systematically follows the crashes. In our simple model, we observe the presence of well defined cycles in the post-crash dynamics, which may vary for different numerical runs, but are identical for all the crashes in the same simulation. The presence of these cycles obviously limits the amplitude of price fluctuations and reduce the volatility. Cycles occur because during the crash, the scores of selling strategies suddenly accumulate a large excess and are thus used for some time after the crash irrespective of their more recent performance. Thus all agents keep using the very same strategies. Since the number of possible past histories is finite (2^m), the dynamics is deterministic and the system must enter a cycle.

Finally, let us mention what happens when tuning the memory β of the scores update. In general we observe in this regime an oscillatory transient before the intermittent behaviour takes over, which lasts approximately 5000 time units. At values of β very close to zero ($\beta \sim 10^{-5}$), we have observed longer transients effects, while already for $\beta \sim 10^{-4}$ we recover the usual dynamical trend. We have also observed that small values of β tend to reduce the number of occasional bubbles and crashes. This can be explained by recalling the mechanism of bubble (crash) formation. As described previously in this section, during the intermittent dynamics bubbles are formed by appropriate statistical fluctuations. However, in order for a bubble to sustain in time, it is crucial that buying strategies emerge as the best performing ones. When β is small the performance of each strategy is evaluated on a very long period and the recent past count relatively little. Thus even the fluctuations due to fundamentalist players may change the performance balance and the bubble cannot self-sustain.

6 Market efficiency and stability

6.1 Dynamical evolution of the parameters and self-organization

As discussed above, our model *a priori* involves a large number of parameters. However, the precise value of most of them is not crucial. The only two important parameters are the ratio g/λ , that measures the impact of trading on prices, and P , that measures the tendency of the agents for trend following ($P > 0$) versus contrarian strategies ($P < 0$). For small values of g/λ , the volatility of price changes is very small, and, as we have discussed in section 4, bubbles form in such a way that very visible trends appear. Furthermore, for small g/λ , these bubbles have a very long lifetime. This means that the agents will naturally increase the fraction of the wealth they invest in this market, since the perceived risk is small. Hence g will spontaneously increase. Simultaneously, small g/λ 's lead to rather small execution rates: see Figure 4. Therefore, the microstructure of the market will evolve such as to make λ smaller, so that the market becomes more liquid. Both effects lead to an increase of the ratio g/λ , and the corresponding destruction of clear trends. This scenario

therefore suggests that g/λ enters the intermittent region, where the market becomes 'quasi-efficient' (*i.e.* returns are uncorrelated on short time scales), but where interesting statistical anomalies appear. The ratio g/λ then stops growing, since the market becomes very volatile, without clear trends. The agents therefore limit their investment. The above scenario is somewhat related to that proposed in the context of the Minority Game in [44]: increasing the number of players leads to a more efficient game, but reduces the incentive for new players to enter, such as a 'marginally efficient' state is reached.

The role of the parameter P is quite interesting. If agents were on average only slightly contrarian, the behaviour of the market would be completely different, with boring small mean reverting fluctuations around the fundamental price, that we have here modeled as a deterministic growth $X_f(t) = \exp(\rho t)$. If this fundamental price was itself randomly fluctuating, our model market would be very close to the standard 'efficient market' picture, where rational agents systematically correct past excess returns such as to lock the market price to the fundamental price. The fact that human psychology seems to favor mimetic, trend following strategies seems to keep the market in the intermittent region of the phase diagram shown in Figure 2.

6.2 Distribution of agents' wealth

We have also studied the distribution of the wealth of the agents in our model. Interesting, for small values of g , the distribution is very uniform across agents: since the level of speculative investment is small, the redistribution effect through trading dominates the random speculative gains. When g increases, however, a Pareto like distribution of wealth, with very unequal agents, sets in, as expected from the analysis of [62].

6.3 The effect of finite transaction costs

An interesting question in the context of financial markets is to understand how the introduction of transaction costs (as the Tobin tax for example) might stabilize the behaviour of these markets. We have studied a generalization of our model in which a non zero proportional fee is paid at each transaction. As we shall see, this drives the system toward a much more stable regime. Transaction costs play a similar role to a negative polarization (more contrarian) strategies.

The equations to be modified are the ones concerning the update of cash, that accounts for the extra cost: equation (12) is changed into:

$$B_i(t+1) = B_i(t)(1 + \rho) - \delta\phi_i(t)X(t+1) - \nu|\delta\phi_i(t)X(t+1)|. \quad (41)$$




Fig. 12. Evolution of the period T of the oscillations as a function of the transaction cost parameter ν , for $g/\lambda = 0.005$ and $g/\lambda = 0.1$. Inset: Evolution of the fraction of active agents as a function of ν for $g/\lambda = 0.75$.

Similarly, we adopt a modified update for the strategies scores (see 13):

$$S_i^\alpha(t+1) = (1 - \beta)S_i^\alpha(t) + \beta(\epsilon_i^\alpha(\mathcal{I}_{t-1})[r(t) - \rho] - \nu|\epsilon_i^\alpha(\mathcal{I}_{t-1})|). \quad (42)$$

Let us first discuss the small g/λ regime, where bubbles appear in the absence of transaction costs. Scores of the buying strategies grow most rapidly at the start of the bubble, where the return $r(t_0)$ is maximum. When the transaction costs are such that $\nu > r(t_0) - \rho$, the bubble will never start. Since the initial growth of the bubble is proportional to g/λ , we expect that bubbles are more robust to costs for larger g/λ 's. If a bubble is formed, is bound to collapse earlier for $\nu > 0$, since the performance of the buying strategy is systematically reduced. Hence the frequency of the oscillations will increase with ν . We show in Figure 12 that this is indeed the case. (Note however that as ν increases, some bubbles do not form, and one observes a mixed behaviour with on and off oscillations.) In the intermittent phase, the introduction of transaction costs favors the inactive strategy, and we observe a decay of the average fraction of active agents as ν increases (see Fig. 12). Large values of transaction costs drive the system toward a more stable regime, that still however exhibits long-range volume fluctuations of the type described in the previous section.

The above discussion suggests that transaction costs would stabilize the markets. However, since the values of g/λ that corresponds to realistic markets is rather large, these costs would have to be substantial (a few percent per trade) in order to really affect the markets. This conclusion could be read backward: a small tax, of the order of a few basis points (*i.e.* 10^{-4}), would probably not change

dramatically the behaviour of markets, and would simply add to already existing costs (brokerage, bid-ask spreads and market impact). Of course, definite statements about, say, the Tobin tax, would need a more realistic model.

7 Summary and conclusion

In this paper, we have presented a detailed study of a rather complex market model, inspired from the Santa Fe artificial market and the Minority Game. Agents have different strategies among which they can choose, according to their relative profitability, with the possibility of *not* participating to the market. The price is updated according to the excess demand, and the wealth of the agents is properly accounted for. The set up of the model involves quite a large number of parameters. Fortunately, only two of them play a significant role: one describes the impact of trading on the price, and the other describes the propensity of agents to be trend following or contrarian.

The main result of our study is the appearance of three different regimes, depending on the value of these two parameters: an oscillating phase with bubbles and crashes, an intermittent phase and a stable 'rational' market phase. The statistics of price changes in the intermittent phase resembles that of real price changes, with small linear correlations, fat tails and long range volatility clustering. We have discussed how the time dependence of these parameters could spontaneously lead the system in the intermittent region. We have analyzed quantitatively the temporal correlation of activity in the intermittent phase, and have shown that the 'random time strategy shift' mechanism proposed in an earlier paper [39] allows to understand these long ranged correlations. Other mechanisms leading to volatility clustering have been reviewed.

We have discussed several interesting issues that our model allows to address, such as the detailed mechanism for bubble formation and crashes, the influence of transaction costs, the distribution of agents wealth, and the role of a limited amount of capital on the long time fluctuations of the price.

Many extensions of the model could be thought of, for example, the diversity of time horizons used by the different agents, the introduction of time dependent fundamental factors, or a larger number of tradable assets. A more interesting path, in our mind, would be to simplify the model sufficiently as to be able to analytically predict the phase diagram of Figure 2, and reach a similar level of understanding as in the Minority Game (for some work in this direction, see [36]).

We thank M. Mézard, who participated to the early stages of this work, for discussions. We have benefited from many conversations with S. Bornholdt, A. Cavagna, D. Challet, D. Farmer, T. Lux, M. Marsili and M. Potters. We thank Thomas Lux for providing some data on the Lux-Marchesi model.

References

1. D.M. Guillaume, M.M. Dacorogna, R.D. Davé, U.A. Müller, R.B. Olsen, O.V. Pictet, *Finance and Stochastics* **1**, 95 (1997)
2. R. Mantegna, H.E. Stanley, *An Introduction to Econophysics* (Cambridge University Press, 1999)
3. J.-P. Bouchaud, M. Potters, *Théorie des Risques Financiers* (Aléa-Saclay, 1997); *Theory of Financial Risks* (Cambridge University Press, 2000)
4. For a review, see J.D. Farmer, *Physicists attempt to scale the ivory towers of finance*, in *Comput. Sci. Engineering*, November 1999, reprinted in *Int. J. Theo. Appl. Fin.* **3**, 311 (2000)
5. A. Lo, *Econometrica* **59**, 1279 (1991)
6. Z. Ding, C.W.J. Granger, R.F. Engle, *J. Empirical Finance* **1**, 83 (1993)
7. R. Cont, M. Potters, J.-P. Bouchaud, *Scaling in stock market data: stable laws and beyond*, in *Scale invariance and beyond*, edited by B. Dubrulle, F. Graner, D. Sornette (EDP Sciences, 1997); M. Potters, R. Cont, J.-P. Bouchaud, *Europhys. Lett.* **41**, 239 (1998)
8. Y. Liu, P. Cizeau, M. Meyer, C.-K. Peng, H.E. Stanley, *Physica A* **245** 437 (1997); P. Cizeau, Y. Liu, M. Meyer, C.-K. Peng, H.E. Stanley, *Physica A* **245**, 441 (1997)
9. A. Arnéodo, J.-F. Muzy, D. Sornette, *Eur. Phys. J. B* **2**, 277 (1998)
10. J.-F. Muzy, J. Delour, E. Bacry, *Eur. Phys. J. B* **17**, 537 (2000); *Physica A* **299**, 284 (2001)
11. S. Ghashghaie, W. Breymann, J. Peinke, P. Talkner, Y. Dodge, *Nature*, **381** 767 (1996)
12. B.B. Mandelbrot, *Fractals and Scaling in Finance* (Springer, 1997)
13. A. Fisher, L. Calvet, B.B. Mandelbrot, *Multifractality of DEM/\$ rates*, Cowles Foundation Discussion Paper 1165
14. J.P. Bouchaud, M. Meyer, M. Potters, *Eur. Phys. J. B* **13**, 595 (1999)
15. J.-P. Bouchaud, A. Matacz, M. Potters, *Phys. Rev. Lett.* **87**, 228701 (2001)
16. C. Hommes, *Quantitative Finance* **1**, 149 (2001)
17. C. Chiarella, X.-Z. He, *Quantitative Finance* **1**, 509 (2001)
18. For an introduction and references, see the very interesting book of R. Schiller, *Irrational Exuberance* (Princeton University Press, 2000)
19. I. Giardina, J.P. Bouchaud, *Proceedings of the Prague Conference on Application of Physics to Economic Modelling, February 2001*, *Physica A* **299**, 28 (2001)
20. A. Beja, M.B. Goldman, *J. Finance* **35**, 235 (1980)
21. R.G. Palmer, W.B. Arthur, J.H. Holland, B. LeBaron, P. Taylor, *Physica D* **75**, 264 (1994), see also B. LeBaron, *Quantitative Finance* **1**, 254 (2001)
22. W.B. Arthur, *Am. Ec. Assoc. Papers Proc.* **84**, 406 (1994)
23. M. Levy, H. Levy, S. Solomon, *Microscopic Simulation of Financial Markets* (Academic Press, San Diego, 2000)
24. D. Challet, Y.-C. Zhang, *Physica A* **246**, 407 (1997)
25. P. Bak, M. Paczuski, M. Shubik, *Physica A* **246**, 430 (1997)
26. J.-P. Bouchaud, R. Cont, *Eur. Phys. J. B* **6**, 543 (1998)
27. J.D. Farmer, *Market Force, Ecology and Evolution*, e-print adap-org/9812005, *Int. J. Theo. Appl. Fin.* **3**, 425 (2000)
28. G. Iori, *Int. J. Mod. Phys. C* **10**, 149 (1999)
29. E. Samanidou, E. Zschischang, D. Stauffer, T. Lux, in *Microscopic models for Economic Dynamics*, edited by F. Schweitzer, *Lecture Notes in Physics* (Springer, Berlin, 2002)
30. D. Challet, M. Marsili, Y.-C. Zhang, *Physica A*, 276, 284 (2000). D. Challet, M. Marsili, R. Zecchina, *Int. J. Theo. Appl. Fin.* **3**, 451 (2000), *Phys. Rev. Lett.* **84** 1824 (2000)
31. T. Lux, M. Marchesi, *Nature* **397**, 498 (1999)
32. T. Lux, M. Marchesi, *Int. J. Theo. Appl. Fin.* **3**, 675 (2000)
33. J.D. Farmer, S. Joshi, *J. Economic Behaviour and Organisation*, **49**, 149 (2002)
34. P. Jefferies, M. Hart, P.M. Hui, N.F. Johnson, *Int. J. Theo. Appl. Fin.* **3**, 443 (2000)
35. P. Jefferies, N.F. Johnson, M. Hart, P.M. Hui, *Eur. Phys. J. B* **20**, 493 (2001)
36. M. Marsili, *Physica A*, **299**, 93 (2001)
37. M. Raberto, S. Cincotti, S.M. Focardi, M. Marchesi, *Physica A* **299**, 320 (2001)
38. J.V. Andersen, D. Sornette, preprint cond-mat/0205423
39. J.P. Bouchaud, I. Giardina, M. Mézard, *Quantitative Finance* **1**, 212 (2001)
40. A. Orléan, *J. Economic Behavior and Organisation* **28**, 257 (1995)
41. J.-P. Bouchaud, R. Cont, *Eur. J. Phys. B* **6**, 543 (1998)
42. J.-P. Bouchaud, *Quantitative Finance* **1**, 105 (2000)
43. A. Cavagna, *Phys. Rev. E* **59**, R3783 (1999); A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, *Phys. Rev. Lett.* **83**, 4429 (1999); J.P. Garrahan, E. Moro, D. Sherrington, *Phys. Rev. E* **62**, R9 (2000)
44. D. Challet, A. Chessa, M. Marsili, Y.C. Zhang, *Quantitative Finance* **1**, 168 (2001)
45. M. Marsili, R. Mulet, F. Ricci-Tersenghi, R. Zecchina, *Phys. Rev. Lett* **87**, 208701 (2001)
46. M. Marsili, D. Challet, *Adv. Complex Systems* **4**, 5 (2001)
47. V. Plerou, P. Gopikrishnan, X. Gabaix, H.E. Stanley *Phys. Rev. E* **66**, 027104 (2002)
48. J.D. Farmer, F. Lillo, R. Mantegna, in preparation
49. J.P. Bouchaud, M. Mézard, M. Potters, in preparation
50. V. Plerou, P. Gopikrishnan, L.A. Amaral, M. Meyer, H.E. Stanley, *Phys. Rev. E* **60**, 6519 (1999)
51. C. Godrèche, J.M. Luck, *J. Stat. Phys.* **104**, 489 (2001)
52. E. Sparre-Andersen, *Math. Scand.* **1**, 263 (1953)
53. U. Frisch, in *Scale invariance and beyond*, edited by B. Dubrulle, F. Graner, D. Sornette (EDP Sciences, 1997)
54. for a review, see J.P. Bouchaud, L. Cugliandolo, J. Kurchan, M. Mézard, *Out of Equilibrium dynamics in spin-glasses and other glassy systems*, in *Spin-glasses and Random Fields*, edited by A.P. Young (World Scientific, 1998), and references therein
55. V. Plerou, P. Gopikrishnan, L.A. Amaral, X. Gabaix, H.E. Stanley, *Phys. Rev. E.* **62**, R3023 (2000)
56. G. Bonnano, F. Lillo, R. Mantegna, *Physica A* **280**, 136 (2000)
57. M. Youssefmir, B. Huberman, *J. Economic Behaviour and Organisation* **32**, 101 (1997)
58. A. Kirman, *Quarterly J. Economics* **108**, 137 (1993), and in preparation
59. S. Bornholdt, *Int. J. Mod. Phys. C* **12**, 667, (2001)
60. S. Alfarano, T. Lux, working paper, WEHIA 2002
61. U.A. Muller, M.M. Dacorogna, R. Davé, R. B. Olsen, O.V. Pictet, *J. Empirical Finance* **4**, 213 (1997)
62. J.P. Bouchaud, M. Mézard, *Physica A* **282**, 536 (2000)