Correlated adaptation of agents in a simple market: a statistical physics perspective

J.P. Garrahan , E. Moro & D. Sherrington

Correlated adaptation of agents in a simple market: a statistical physics perspective

J P Garrahan, E Moro and D Sherrington

Theoretical Physics, University of Oxford, 1 Keble Road, Oxford
OX1 3NP, UK

Abstract

We discuss recent work in the study of a simple model for the collective behaviour of diverse speculative agents in an idealized stockmarket, considered from the perspective of the statistical physics of many-body systems. The only information about other agents available to any one is the total trade at time steps. Evidence is presented for correlated adaptation and phase transitions/crossovers in the global volatility of the system as a function of appropriate information scaling dimension. Stochastically controlled irrationality of individual agents is shown to be globally advantageous. We describe the derivation of the underlying effective stochastic differential equations which govern the dynamics, and make an interpretation of the results from the point of view of the statistical physics of disordered systems.

1. Introduction

There is currently much interest in the physics community in complex cooperative behaviour of systems of many individual entities influencing one another competitively. In particular, when combined with non-uniformity in the inclinations of the individuals, the behaviour of the whole can exhibit much greater complexity, richness and subtlety than is present in the rules governing individuals; in the words of P W Anderson 'more is different' [1]. Examples are found in spin glasses (disordered magnetic alloys), neural networks and hard optimization problems (for recent reviews see [2, 3]). Economic markets also involve many individuals whose desires are not all simultaneously satisfiable and who often have different inclinations and strategies. It is therefore natural to ask to what extent the problems in physics and in economics are similar or different, to what extent the techniques and concepts developed for the physics problems can be applied to those in economics, to what extent the economics problems pose new challenges for the physicists, and to search for fruitful symbiosis of understanding, quantification and application. To this end we discuss in this paper recent developments in the study of a model inspired by economics, but analysed from the perspective of physics, finding unexpected new results and

subtleties, and concluding that both subjects have something to teach each other and that there is potential for further transfers and discoveries. We make no attempt to be encyclopaedic or chronologically historical.

Before discussing the specific model, some general remarks are appropriate. Much of the progress in physics has come from starting with the simplest but non-trivial microscopic entities and interaction rules which can still lead to complex behaviour at the macroscopic[1] level. Greater 'reality' at the microscopic level can be added later. Many results are robust to the microscopic details, although new features can also arise with sufficient qualitative change. We apply a similar philosophy here, deliberately oversimplifying at the individual level to expose novel consequences of cooperation uncluttered by microscopic complication. Thus, we concentrate on systems with simple microscopic dynamical rules and minimal number of control parameters. In the spirit of statistical many-body theory we concentrate on systems with many ($N \gg 1$) microscopic players, with particular regard to the leading large-N behaviour of macroscopic quantities. We allow for temporally-fixed variation among individuals but,

---

[1] By 'macroscopic' we refer to quantities which are averaged over the behaviour of all the 'microscopic' individuals.

in the spirit of statistical relevance, we draw their individual characteristics independently from identical distributions. We also allow for stochasticity (temporal indeterminacy) at the individual operational level, but again in a statistically relevant and minimal parameter fashion. As usual in statistical physics, we expect self-averaging[2] of normal macroscopic observables, although non-self-averaging might be envisaged at a more sophisticated level [3, 4].

It is also appropriate to contrast our study with those of conventional economics theory. A typical assumption used in neoclassical economic theory [5]—especially game theory [6]—is that agents are hyper-rational. They know the utility functions of other agents, they are fully aware of the process they are embedded in, they make optimum long-run plans, and so forth. This is a rather extravagant and implausible model of human behaviour, especially in situations like a stock market. Moreover, in neoclassical economic theory, microscopic equilibrium is the reigning paradigm [5]. Individual strategies are assumed to be optimal given expectations, and expectations are assumed to be justified given the evidence. Equilibrium is thus reached in one-step dynamics once hyper-rationality is assumed. In this paper we consider a different, perhaps more realistic, scenario in which the only information available to any agent about the others is of the macroscopic consequences of the multiplicity of their actions (i.e. the analogues of market indices). We allow for diversity and irrationality in that they do not all draw the same conclusions from this information [7], nor do they necessarily operate deterministically. In general there will not be microscopic equilibrium although there may be macroscopic equilibrium[3].

The paper is organized as follows. In section 2 we present the Minority Game [8] and review its main features. This model is a specific realization of Arthur's 'El Farol' Bar Problem [7], and is the starting point of our investigations. In section 3 we consider a continuous generalization of the model, and study the effect of allowing for stochastic decision making on the part of the agents. The derivation of a fundamental analytic theory is discussed in section 4, where the underlying stochastic differential equations for the dynamics of the system are presented. Section 5 contains an interpretation of the results and we give our conclusions in section 6.

2. The Minority Game

The model system we consider is one known as the Minority Game (MG) [8] and is intended to mimic in a simplified way a market of agents bidding to profit by buying when the majority wish to sell (so that the price can be lowered) and selling when the majority wish to buy (so that a higher price can be negotiated) [7]. It comprises a large number of agents each of whom can act as buyer or seller, deciding on how to play at each time step through the application of a personal strategy to commonly available information[4]. Each agent has a small set of available strategies, drawn randomly, independently and immutably with identical probabilities from a large suite of strategies. At each time step each agent picks one of his or her strategies, based on points allocated cumulatively to the strategies according to their (virtual) performance in predicting the minority action. For simplicity no other rewards are given.

The system has quenched (fixed in time) randomness in the set of strategies picked at the start of the game by each agent, and it has frustration[5] in that the rewards are for minority action, so not all the individual inclinations can be satisfied simultaneously. There is no direct interaction between agents. Nevertheless, correlation does arise through the adaptive evolution of the use of strategies and manifests itself in the interesting non-trivial macroscopic behaviour of the system.

In the original formulation of the MG [8] agents could only make two choices, buy or sell, with no weight attributed to the size of the order. Conventionally, the strategy points are set initially to zero and thus the agents start making their first choice at random. The common information on which they based their decisions was the minority choice (buy or sell) over the last m time steps. The strategies were quenched randomly chosen Boolean functions acting on this information, the binary output determining the buy/sell decision. The strategy used by any agent at any time was the one with the currently greatest point-score from those at his/her disposal. Numerical simulations showed [9] that while the average in time (and over the realizations of the disorder) of the total action was just an equality of buyers and sellers, due to the symmetric nature of the model, the standard deviation of its fluctuations away from this value (the analogue of the volatility of a conventional stockmarket) displayed remarkable structure. As a function of the 'memory' m, the volatility has two regimes: for low values of m, the volatility is larger than the value corresponding to all agents just playing randomly; it decreases monotonically with increasing m, crosses the value corresponding to random behaviour, and reaches a minimum at a critical value of the memory m_c; it then starts to grow monotonically with m, asymptotically approaching the random value from below. This non-trivial behaviour of the fluctuations was interpreted as an indication of a cooperative 'phase transition' in the system [9, 10] (see, e.g., [11], for a introduction to phase transitions). Simulations also showed [9] that the relevant scaling variable was the reduced dimension of the space of strategies $d = 2^m/N$, and that the volatility scaled with the number of agents as $\sqrt{N}$.

A second interesting observation was made in [12], where it was shown numerically that the macroscopic behaviour of the MG was unaffected by replacing the time history by

---

[2] By 'self-averaging' we mean that, in the limit $N \to \infty$, the value in a typical realization is the same as the average over realizations.

[3] By 'microscopic equilibrium' we mean a situation in which it is disadvantageous for any individual to change his state, whereas 'macroscopic equilibrium' refers to a situation in which the thermodynamically relevant (leading N) macroscopic observables do not change with time, even though individual microscopic states do change.

[4] A strategy is an operator which acts on a set of data, referred to as the 'information', to yield an outcome which is a buy or sell instruction.

[5] 'Frustration' refers to an inability to satisfy simultaneously all the inclinations of all the microscopic entities, and is believed to be a fundamental ingredient in producing the complexity observed in glassy systems (where quenched disorder is often also present) [2–4].
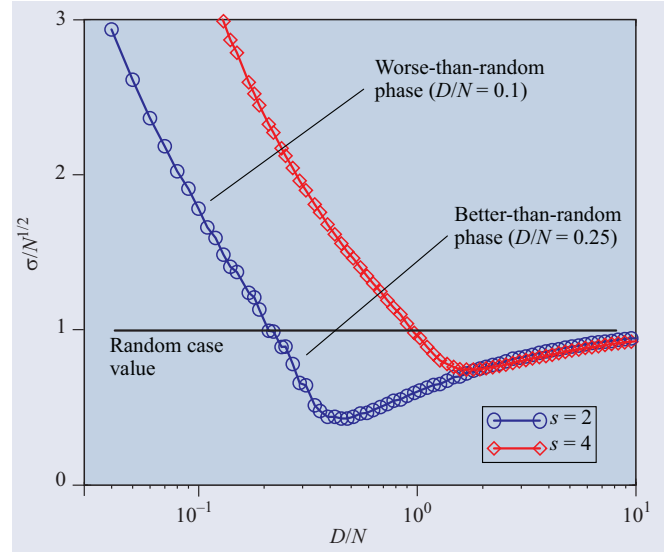
an artificial history, chosen randomly and independently at each time step from the space of all possible histories with uniform probability, provided all agents received the same bogus information (and that the point-scores were still updated on the basis of performance). This is a consequence of the ergodic nature of the MG [13], and indicates that the principal role of the information is in providing a correlation mechanism between the agents in terms of the strategies used. This observation offers a great simplification for the analysis of the model, since replacing the true history by just external noise allows us to study a simpler system which is stochastic but local in time, instead of the more difficult original problem which is deterministic but non-local in time (see also [14]). In fact, we shall also see the consequence of allowing for a further different kind of stochasticity in the next section.

3. The thermal Minority Game

As discussed in the introduction, it is natural to expect that the qualitative features of the MG are robust under changes in the microscopic detail of the model. In this section we present a generalization of the MG to continuous degrees of freedom and to allow for stochastic decision making on the part of the agents, known as the thermal Minority Game (TMG), and first introduced in [13]. This generalization not only preserves the main features of the MG discussed so far, but also gives interesting and advantageous new behaviour, and enables a simpler analytic description of the coarse-grained microdynamics of the system.

3.1. Continuous MG

The continuous version of the MG is as follows [13]. The system consists again of N agents playing the game. At each time step t, each agent reacts to a common piece of 'information' $\vec{\mathcal{I}}(t)$, by taking an action or bid $b_i(t)$ ($i = 1, \ldots, N$). Following the observation of [12] the information $\vec{\mathcal{I}}(t)$ is taken to be a random noise, defined as a unit-length vector in a D-dimensional space[6], for instance $\mathbb{R}^D$, δ-correlated in time and uniformly distributed on the unit sphere[7]. The bid $b(t)$ is defined to be a real number, which can be interpreted as placing an order in a market, of size $|b(t)|$ and positive/negative meaning buy/sell. Bids are prescribed by strategies: maps from information to bid, $\mathbb{R}^D \to \mathbb{R}$. For simplicity the strategy space Γ of the model is restricted to the subspace of homogeneous linear mappings, in contrast to the whole space of binary functions in the MG. Thus a strategy $\vec{R}$ is defined as a vector in $\mathbb{R}^D$, subject for normalization to the constraint $\|\vec{R}\| = \sqrt{D}$, and the prescribed bid is the scalar product $\vec{R} \cdot \vec{\mathcal{I}}(t)$. Each agent has S strategies, drawn randomly and independently from Γ with uniform distribution, remaining fixed throughout the game. In what follows we will focus for simplicity on the case of two strategies per agent, $S = 2$, the generalization to $S > 2$ being straightforward (and

---

[6] Here D is the analogue of 2^m in the original model.
[7] This is a convenient choice. Any other normalized isotropic distribution in $\mathbb{R}^D$, e.g. a Gaussian, would be qualitatively equally suitable. The same applies to the strategies.

Figure 1. Continuous formulation of the MG [13]: the scaled volatility $\sigma/\sqrt{N}$ as a function of the reduced dimension D/N, at $S = 2$ and $S = 4$. The horizontal line is the variance in the random case. The total time t and the initial time t_0 are 10 000 steps. Average over 100 samples, $N = 100$.

the case $S = 1$ being trivial, since there is no possibility of adaptation on the part of the agents). At time step t each agent i chooses one of his/her strategies $\vec{R}_i^\star(t)$ to play with. The 'total bid' (or 'excess demand') is then $A(t) \equiv \sum_j b_j(t) = \sum_j \vec{R}_j^\star(t) \cdot \vec{\mathcal{I}}(t)$. The agents keep track of the potential success of the strategies by assigning points to them, which are updated according to $P(\vec{R}, t + 1) = P(\vec{R}, t) - A(t) b(\vec{R})/N$, where $P(\vec{R}, t)$ represents the points of strategy $\vec{R}$ at time t.

Let us now see whether the results obtained with this continuous formulation of the MG are the same as in the original binary setup. To this end we first review the results of simulations. The average of the total bid $A(t)$ over time and quenched disorder is zero, as is expected from the symmetric nature of the model. As discussed in the previous section, the first relevant macroscopic observable is the standard deviation σ of the total bid, or volatility, $\sigma^2 \equiv N^{-1}\overline{\langle A^2(t)\rangle}$, where the overline means disorder average, $\langle \cdot \rangle \equiv \lim_{t \to \infty} \frac{1}{t} \int_{t_0}^{t_0+t} (\cdot)\, dt'$, and we have normalized σ by $\sqrt{N}$ according to the findings of [9]. In figure 1 we show that the main features of the MG are reproduced: first, the relevant scaling parameter is the reduced dimension of the strategy space $d = D/N$; second, the volatility starts for low d at a value larger than the one corresponding to the agents choosing randomly, $\sigma_r = 1$ in this case, decreases monotonically until it reaches a minimum at $d = d_c(S)$, the minimum being shallower the higher is the number of strategies S [15], and then it approaches σ_r asymptotically from below. It is easy to check that *all* the other standard features of the binary model are reproduced in the continuous formulation.

3.2. Stochastic decision making

In the original formulation of the MG the agents played in a deterministic fashion using their 'best' strategies, the ones with

the highest number of points. In this subsection we introduce indeterminacy (irrationality) on the part of the agents and show that it can be advantageous.

In the TMG a natural generalization to non-deterministic behaviour is allowed [13]. At time step t, each agent i chooses $\vec{R}_i^{\star}(t)$ randomly from his/her $\{\vec{R}_i^a\}$ with probabilities $\{\pi_i^a(t)\}$ ($a = 1, \ldots, S$). The probabilities $\pi_i^a(t)$ are functions of the points parametrized by a 'temperature' T, defined so as to interpolate between the MG case at $T = 0$, all the way up to the totally random case $\pi_i^a = 1/S$ at $T = \infty$. The temperature can be thought of as a measure of the power of resolution of the agents: when $T = 0$ they are perfectly able to distinguish which are their best strategies, while for increasing T they are more and more confused, until for $T = \infty$ they choose their strategy completely at random. In the language of Game Theory, when $T = 0$ agents play 'pure' strategies, while at $T > 0$ they play 'mixed' ones [6].

We now consider the consequences of having introduced the temperature. For simplicity we assume $S = 2$. For the probabilities we choose the form $\pi_i^{1,2}(t) \propto \exp[\pm\beta\,\mathrm{sgn}(p_i(t))]$ (with $\sum_a \pi_i^a(t) = 1$ and $\beta = 1/T$) [19], which satisfies the requirements of the previous paragraph. Consider now a value of d belonging to the worse-than-random region of the MG (see figure 1) and let us see what happens to the volatility σ when we switch on the temperature. We know that for $T = 0$ we must recover the same value as in the ordinary MG, while for $T \to \infty$ we expect to obtain the value σ_r of random choice. But in between a very interesting thing occurs: $\sigma(T)$ is not a monotonically decreasing function of T, but there is a large intermediate temperature regime where σ is *smaller* than the random value σ_r; see figure 2. The meaning of this result is the following: even if the system is in a MG phase which is worse than random, there is a way to significantly decrease the volatility σ below the random value σ_r by *not* always using the best strategy, but rather allowing a certain degree of individual error.

Furthermore, even if we fix d at a value belonging to the better-than-random region, but with $d < d_c$, a similar range of temperature still improves the behaviour of the system, decreasing the volatility even below the MG value (see figure 2). In the phase $d > d_c$ the behaviour changes; the optimal value of σ is at $T = 0$, and the volatility simply increases with increasing temperature towards σ_r, as shown in figure 2.

Another possible functional form for the probabilities is $\pi_i^a(t) \propto \exp[\beta P(\vec{R}_i^a, t)]$. For long but finite simulation time this yields the results shown in figure 3(a) [13], which are analogous to those of figure 2. However, the upturn of σ for large temperatures for this choice of probabilities is only a transient [16–18]: for any $T \gg 1$, if one waits for times of order NT the volatility stays at its smallest possible value; see figure 3(b). Thus, we see that in the $d < d_c$ phase, for any finite temperature the performance of the system will be better than the original MG, and for *any finite* temperature $T > \mathcal{O}(1)$ the performance will be optimal, provided one waits long enough.
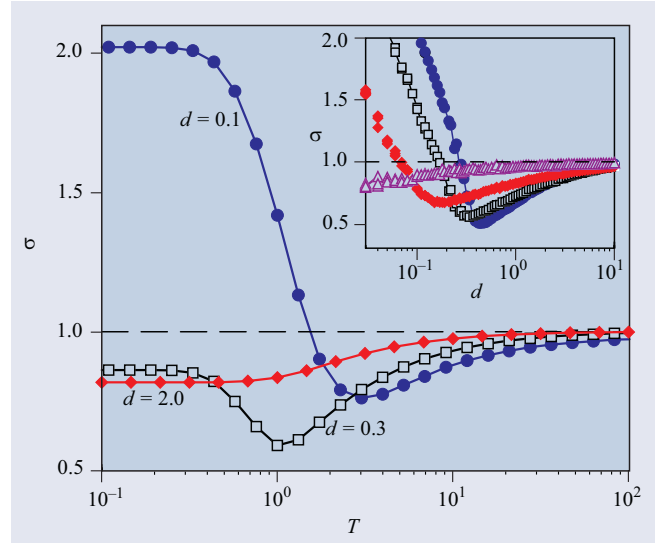
Figure 2. TMG: volatility as a function of the temperature for the case of probabilities $\pi_i^{1,2}(t) \propto \exp[\pm\beta\,\mathrm{sgn}(p_i(t))]$. Inset: volatility as a function of d for different values of the temperature $T = 10^{-3}, 1, 2, 10$.

4. Continuous time dynamics of the TMG

In this section we derive the continuous time limit of the TMG as a starting point for the analytical study of its dynamics [19]. We do this in two steps. We first show that, to a good approximation, the external information can be eliminated in favour of an effective interaction between the agents. We then cast the dynamical equations of the TMG as a set of stochastic differential equations for an interacting disordered system with non-trivial random diffusion. Again, for simplicity, we consider explicitly $S = 2$.

The set of unconstrained degrees of freedom of the TMG is given by the differences $\{p_i(t)\}$ of the points of the two strategies of each agent. The choice of strategies used at each stage is given by $\vec{R}_i^{\star}(t) = \vec{\omega}_i + \vec{\xi}_i\,\mathrm{sgn}[s_i(t) + \mu(t)]$, where $\vec{\omega}_i \equiv \left(\vec{R}_i^1 + \vec{R}_i^2\right)/2$, $\vec{\xi}_i \equiv \left(\vec{R}_i^1 - \vec{R}_i^2\right)/2$, $s_i(t) \equiv \pi_i^1(t) - \pi_i^2(t)$, and $\mu(t)$ is a stochastic random variable uniformly distributed between -1 and 1 and independently distributed in time. The equations for the point differences then read

$$p_i(t+1) = p_i(t) - (\vec{a}(t) \cdot \vec{\mathcal{I}}(t))\,(\vec{\xi}_i \cdot \vec{\mathcal{I}}(t)), \qquad (1)$$

where $\vec{a}(t) \equiv \sum_i \vec{R}_i^{\star}(t)/N$. Equation (1), together with the random processes for $\vec{\mathcal{I}}(t)$ and $\vec{R}_i^{\star}(t)$, define the dynamics of the TMG.

We now consider the continuous time limit of equation (1) in such a way as to preserve all the macroscopic features of the TMG. To this end we introduce an arbitrary time step Δt. We deal first with the information $\vec{\mathcal{I}}(t)$. Let us assume that $\vec{\mathcal{I}}(t)$ is a differential random motion in the space of strategies, i.e., $\vec{\mathcal{I}}(t) = \Delta\vec{W}(t)$, with zero mean and variance Δt. Replacing in equation (1) we obtain $p_i(t + \Delta t) = p_i(t) - \vec{a}(t) \cdot \Delta\vec{W}(t)\,\vec{\xi}_i \cdot \Delta\vec{W}(t)$. In the limit $\Delta t \to 0$, and using the Kramers–Moyal
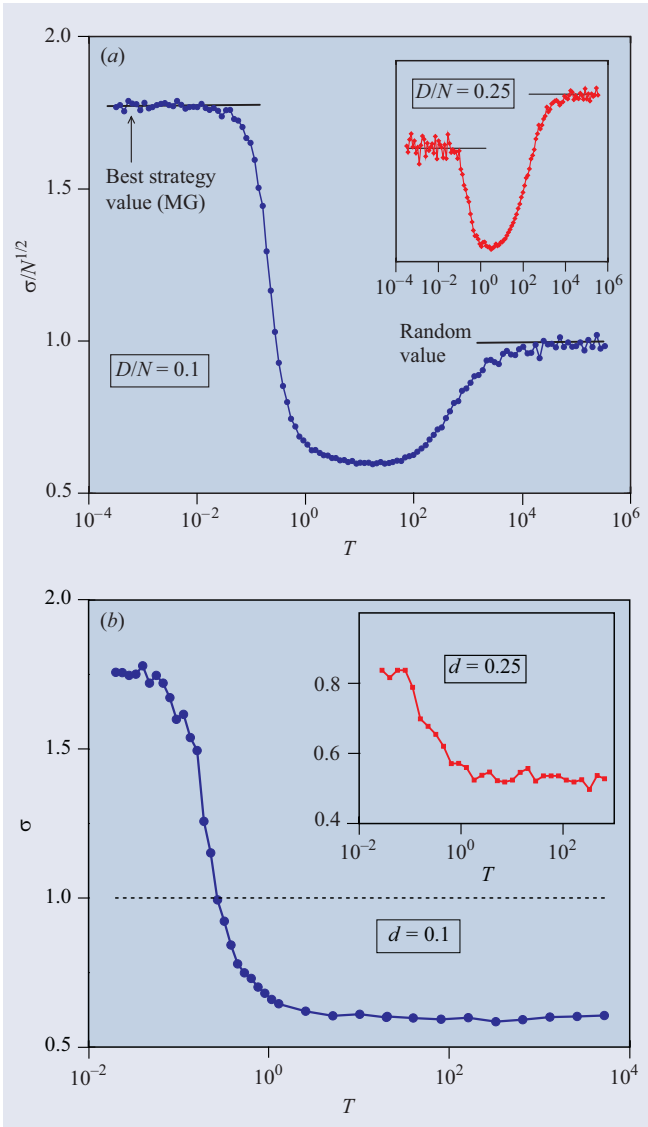
Figure 3. (*a*) TMG: volatility as a function of the temperature T at $d = 0.1$, for the case $\pi_i^a(t) \propto \exp[\beta P(\vec{R}_i^a, t)]$, for finite waiting time. In the inset we show $d = 0.25$. (*b*) The same as in (*a*) for waiting times larger than NT.

expansion [20], we get

$$\mathrm{d}p_i(t) = -\frac{1}{ND} \sum_j \vec{R}_j^\star(t) \cdot \vec{\xi}_i \, \mathrm{d}t + \mathcal{O}(\mathrm{d}t^2). \quad (2)$$

Note that to $\mathcal{O}(dt)$ the noise has been eliminated in favour of an effective strategy interaction among the agents, and the standard deviation becomes

$$\sigma^2 = \frac{1}{ND} \overline{\left\langle \sum_{ij} \vec{R}_i^\star(t) \cdot \vec{R}_j^\star(t) \right\rangle}. \quad (3)$$

When the temperature is different from zero the TMG equation (2) still depends on the stochastic choice of strategies $\vec{R}_i^\star(t)$, even at leading order. At each time step, $\vec{R}_i^\star$ takes one of the two possible values $\vec{R}_i^{1,2}$, defining a stochastic jump process. In order to write the corresponding master

equation we need to know the transition probabilities. The rhs of equation (2), which we denote Δ_i, is a normalized sum of N random numbers $\vec{\xi}_i \cdot \vec{R}_j^\star(t)$, each with mean $m_{ij} = \vec{\xi}_i \cdot \vec{\omega}_j + (\vec{\xi}_i \cdot \vec{\xi}_j) s_j(t)$, and variance $v_{ij} = (\vec{\xi}_i \cdot \vec{\xi}_j)^2 [1 - s_j^2(t)]$. By the central limit theorem, we know that for N large Δ_i will tend to be normally distributed with mean $\langle\langle \Delta_i \rangle\rangle = \sum_j m_{ij}$, and variance $\langle \Delta_i^2 \rangle = \sum_j v_{ij}$, where $\langle\langle \cdot \rangle\rangle$ stands for the average over realizations of the random process $\mu(t)$. Moreover, Δ_i and $\Delta_{j\neq i}$ are correlated, the covariance matrix given by

$$M_{ij}[\boldsymbol{p}(t)] \equiv \langle\langle \Delta_i \Delta_j \rangle\rangle - \langle\langle \Delta_i \rangle\rangle \langle\langle \Delta_j \rangle\rangle$$
$$= \sum_k (\vec{\xi}_i \cdot \vec{\xi}_k)(\vec{\xi}_j \cdot \vec{\xi}_k) \left[1 - s_k^2(t) \right], \quad (4)$$

where $\boldsymbol{p} \equiv (p_1, \ldots, p_N)$, etc. Collecting these results, we obtain the transition probabilities in the large-N limit, $W(\boldsymbol{p}'|\boldsymbol{p}) = \Phi(\nabla_s \mathcal{H}; \mathbb{M})$, where Φ corresponds to the normal distribution with mean $\nabla_s \mathcal{H}$ and covariance matrix $\mathbb{M} \equiv \{M_{ij}\}$, where the 'Hamiltonian' $\mathcal{H}$ is given by[8]

$$\mathcal{H} = \frac{1}{2}\Omega + \sum_i h_i s_i + \frac{1}{2} \sum_{ij} J_{ij} s_i s_j, \quad (5)$$

with

$$h_i \equiv \sum_j \vec{\omega}_j \cdot \vec{\xi}_i / ND, \quad J_{ij} \equiv \vec{\xi}_j \cdot \vec{\xi}_i / ND. \quad (6)$$

Note that $\partial \mathcal{H}/\partial s_i \sim \mathcal{O}(1)$, and $M_{ij} \sim \mathcal{O}(1/N)$, so that fluctuations are also of $\mathcal{O}(1)$ and thus are *not* suppressed when $N \to \infty$.

The $\mu(t)$ are chosen independently at each time. If we make the natural assumption that in the limit $dt \to 0$ their correlation at different times is a δ-function, the master equation becomes a Fokker–Planck equation by means of Kramers–Moyal expansion [20]

$$\frac{\partial \mathcal{P}}{\partial t} = -\sum_i \frac{\partial}{\partial p_i} \left(\frac{\partial \mathcal{H}}{\partial s_i} \mathcal{P} \right) + \frac{1}{2} \sum_{ij} \frac{\partial^2}{\partial p_i \partial p_j} \left(M_{ij} \mathcal{P} \right). \quad (7)$$

The dynamics of the TMG is therefore effectively described by a set of stochastic differential equations for the point differences

$$\mathrm{d}\boldsymbol{p} = -\nabla_s \mathcal{H} \, \mathrm{d}t + \mathbb{M} \cdot \mathrm{d}\boldsymbol{W}, \quad (8)$$

where $\boldsymbol{W}(t)$ is an N-dimensional Wiener process, and the volatility is given by $\sigma^2 = 2\overline{\langle \mathcal{H} \rangle} + \sum_i \overline{J_{ii}} - \sum_i \overline{J_{ii} \langle s_i^2 \rangle}$.

A detailed interpretation of $\mathcal{H}$ is given in the next section, but briefly equation (8) is suggestive of it as a controlling energy with the 'motion' of $\boldsymbol{p}$ given by its derivative. However we note that for a natural analogue of Newton's law or its generalization to a noisy environment, as in the Langevin equation, the derivative would be with respect to $\boldsymbol{p}$, whereas here it is with respect to $\boldsymbol{s}$ (which is a function of $\boldsymbol{p}$), so that a metric is needed to relate $\mathrm{d}\boldsymbol{p}/\mathrm{d}t$ to the natural force $\nabla_s \mathcal{H}$. At finite temperature one also has the unusual extra diffusive/noise term $\mathbb{M} \cdot \mathrm{d}\boldsymbol{W}$. An investigation based on replacing $\nabla_s \mathcal{H}$ by $\nabla_p \mathcal{H}$ and ignoring the diffusive term was performed in [26].

---

[8] This expression was first obtained, by a different procedure and with a different interpretation, in [10].
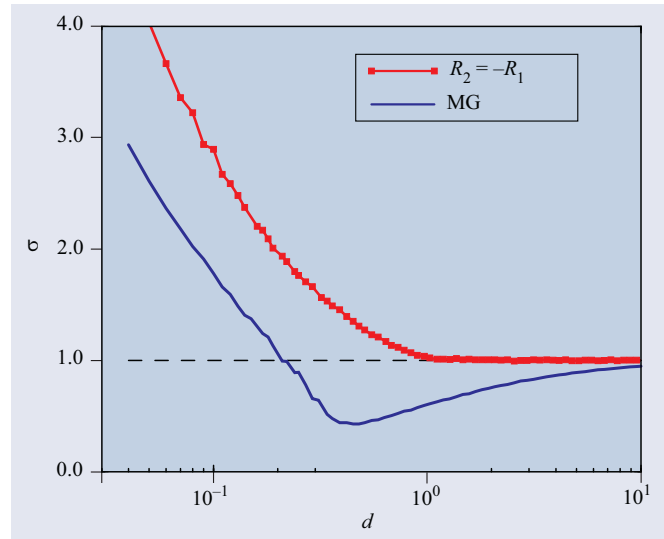
5. Interpretations

Notwithstanding the subtleties mentioned in the last paragraphs, it is interesting to consider the implications of $\mathcal{H}$ as a controlling function of the dynamics. The form of $\mathcal{H}$ exhibited in equation (5) is a familiar one in statistical physics, with $-J_{ij}$ interpreted as measuring the strength of correlation between spins s_i and s_j and $-h_i$ as a (magnetic) field acting on s_i and trying to 'orient' it. Hence is formally justified the interpretation of common 'information', to which all respond, as providing a mechanism of effective mutual interaction[9]. More particularly, taking account of the random character of $\{h_i\}$ and $\{J_{ij}\}$, $\mathcal{H}$ is reminiscent of the Sherrington–Kirkpatrick model [24] of a spin glass and the Hopfield model [25] of a neural network, in both cases augmented by random fields [2]. In the SK model the J_{ij} are chosen randomly from a distribution with variance scaling as N^{-1}, while for the Hopfield model J_{ij} is as given in equation (6), but with opposite sign and where the Cartesian components of $\{\vec{\xi}_i\}$ correspond to memorized patterns of activity. In both of these models (SK and Hopfield) the ground-state energy (minimum of the corresponding $\mathcal{H}$) is less than the energy associated with a paramagnet (value of $\mathcal{H}$ corresponding to random $\{s_i\}$) due to judicious correlation of the $\{J_{ij}\}$ and $\{s_i\}$, even with all the $\{h_i\} = 0$. It is then natural to ask whether a similar correlation is responsible for the reduction of the volatility of the MG below that for random operation. This is however not the case: in figure 4 [32] we show the result of numerical simulations of a system in which the second strategy of each agent $\vec{R}_i^2$ is exactly the opposite of its first strategy, that is $\vec{R}_i^2 = -\vec{R}_i^1$ (with $\vec{R}_i^1$ still chosen randomly at the start of the game), so that all the h_i are zero, and we note that σ never falls below the random value. Clearly, the random field term is crucial in reducing the volatility advantageously [10, 31].

The first term in the dynamical equations (8) corresponds (up to a factor) to gradient descent in the surface defined by $\mathcal{H}$. The natural question then is whether the asymptotic states reached by the dynamics are given by the typical extrema of $\mathcal{H}$, which would imply that in the long run the system reaches macroscopic equilibrium. This issue was explored in [26] (see also [27]), where $\mathcal{H}$ was minimized employing techniques developed for the statistical mechanics of spin glasses [24, 28] and adapted for neural networks [29] (for reviews see [3,4,30]). An excellent agreement with the numerics of the MG was found in the phase $d > d_c$, but this method failed to reproduce the behaviour in the $d < d_c$ phase (see figure 5), although an equilibrium transition was found at a value of d close to d_c[10].

This seems to suggest that the behaviour of the MG in the $d < d_c$ phase is dynamical in nature. A simple test is the

Figure 4. Volatility for anticorrelated strategies $\vec{R}_i^2 = -\vec{R}_i^1$ (squares) compared with the standard MG (solid line).

sensitivity to initial conditions. In figure 5 we show the results of simulating the original dynamics equation (1) of the MG starting from any initial conditions with $|p_i(0)| \sim \mathcal{O}(1)$ [19], instead of the conventional choice of $p_i(0) = 0$[11]. We can see that the the behaviour of the system is very different in the region $d < d_c$: after an initial transient, the variance falls below the initial random value and stays in the better-than-random phase for all values of d. This sensitivity of the results to the initial conditions is a clear indication that the system does not equilibrate for $d < d_c$. An important open problem in the MG is finding the relation between the equilibrium (static) phase transition obtained from minimization of the Hamiltonian $\mathcal{H}$ [26] and the actual dynamical transition observed in the simulations[12].

The differential equations (7) and (8) provide the basis for development of a macroscopic dynamics, either via the dynamical replica theory of Coolen and Sherrington [33] or via the generating functional approach [34]. However, we defer this discussion to later papers.

Finally we remark on the relationship with the crowd–anticrowd concept of [35], where a crowd is a group of agents playing the same strategy and the corresponding anticrowd play the opposite one. The proposal of [35] is that the macroscopic properties of the MG can be described by the behaviour of the crowd–anticrowd pairs. In this approach the volatility of the system is approximated as the sum over all the pairs of crowd–anticrowds of the square of the difference of their sizes. We can formalize this concept in the continuous formulation of the MG by defining the crowd–anticrowd pairs in terms of the projection of the strategies used by the agents on an arbitrary orthonormal basis $\{\hat{e}_\mu\}$ ($\mu = 1, \ldots, D$) of

---

[9] In fact, in conventional statistical physics, the converse is often employed in formal analysis, replacing direct interaction by interaction through randomly distributed global intermediates [21–23].

[10] In the case of the TMG with probabilities $\pi_i^a(t) \propto \exp[\beta P(\vec{R}_i^a, t)]$, the asymptotic value of the volatility for large T (see figure 3(b)) coincides with the one predicted from the equilibrium calculation of [26] (see also [17]). This can be understood from equation (8): for large T, a systematic rescaling of time and points can be used to eliminate the diffusive term in the equations, and the effective dynamics becomes independent of T [19]. Note that this is a consequence of the functional form of the probabilities, and does not hold in the case $\pi_i^{1,2}(t) \propto \exp[\pm\beta \, \text{sgn}(p_i(t))]$.

[11] Note that with the conventional choice the agents do not prefer initially any of their strategies, while when $|p_i(0)| \sim \mathcal{O}(1)$ agents have initially a preferred strategy.

[12] In the case of anticorrelated strategies of figure 4, minimization of $\mathcal{H}$ gives an equilibrium transition at $d = 1$ [31] while the $d < 1$ phase is out-of-equilibrium.
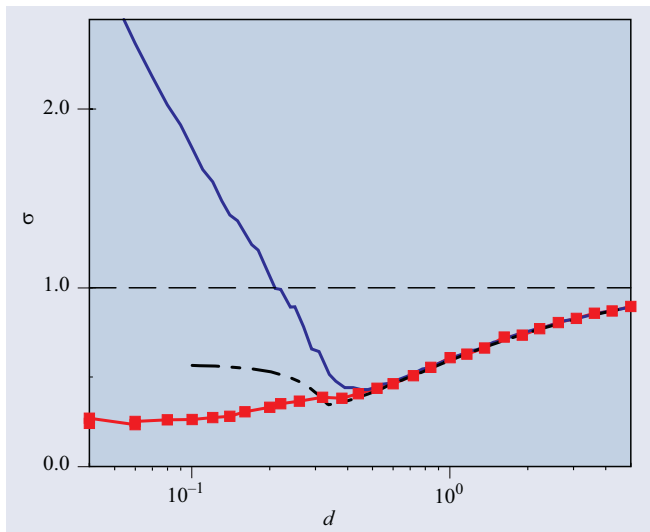
Figure 5. Volatility as a function of d for random initial conditions (squares) compared with the standard MG (solid line). The dotted line corresponds to the approximation of [26].

the strategy space: $N_\mu - N_{\bar{\mu}} \equiv \sum_i \hat{e}_\mu \cdot \vec{R}_i^\star$. This definition is analogous to that of 'staggered' magnetizations in a spin system or a neural network [29]. If we write the volatility using the approximation of [35] $\sigma^2 = \frac{1}{N} \sum_\mu (N_\mu - N_{\bar{\mu}})^2$, and make use of the completeness relation of the $\{\hat{e}_\mu\}$, we recover equation (3). Moreover, an effective dynamics of crowds–anticrowds, as proposed in [36], is exactly derivable from the microscopic dynamical equations (8).

It is of course interesting to ask about the stability of our conclusions to minor perturbations of the model. We have already noted the stability of the qualitative features of the original Minority Game to a change in the number of strategies per agent (excluding the special case $S = 1$). This continues to the thermal extension. We have not investigated explicitly changes to the learning rule but it seems reasonable to expect qualitatively analogous behaviour for other generalized minority rules, for example nonlinear rewards but still favouring minority action [37]. Reward rules involving capital accumulation and consequent variation of potential market influence [38] are another natural extension, as also evolution of the strategies themselves [39, 40]

6. Conclusion

In this paper we have discussed the application of techniques and philosophy of the statistical physics of complex cooperative frustrated many-body systems to a simple model of agents in a competitive market. We have shown that this can lead to both qualitatively and quantitatively new results in the economics-inspired model and also, conversely, that economics models can yield new challenges for statistical physics. The techniques have included computer simulation and analysis. In the simulations we have concentrated on the macroscopic steady-state. In the analysis we have derived a potentially useful microscopic formulation in terms of stochastic differential equations, itself different and

more subtle than that normally encountered in conventional condensed-matter statistical physics at the corresponding coarse-grained microscopic level. The challenge still remains to develop a full macrodynamics, both equilibrium and out-of-equilibrium, but the derived microdynamics is the necessary ingredient for extension of relevant techniques from statistical physics.

We have restricted discussion to a model which is clearly oversimplified from the perspective of a true economic market, but it is possible to envisage still simplified but more realistic models which should be capable of more truly analysing the meaning of 'efficiency' and going beyond it.

Acknowledgments

References

[1] Anderson P W 1972 More is different—broken symmetry and nature of hierarchical structure in science *Science* **177** 393
[2] Young A P (ed) 1998 *Spin Glasses and Random Fields* (Singapore: World Scientific)
[3] Sherrington D 1999 Spin glasses *Physics of Novel Materials* ed M P Das (Singapore: World Scientific) p 146
[4] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
[5] See, e.g., Kreps D M 1990 *A Course in Microeconomics Theory* (New York: Harvester Wheatsheaf)
[6] Fudenberg D and Tirole J 1991 *Game Theory* (Cambridge, MA: MIT Press)
[7] Arthur W B 1994 Inductive reasoning and bounded rationality *Am. Econ. Rev.* **84** 406
[8] Challet D and Zhang Y C 1997 Emergence of cooperation and organization in an evolutionary game *Physica* A **246** 407
[9] Savit R, Manuca R and Riolo R 1999 Adaptive competition, market efficiency and phase transitions *Phys. Rev. Lett.* **82** 2203
[10] Challet D and Marsili M 1999 Phase transition and symmetry breaking in the Minority Game *Phys. Rev.* E **60** R6271
[11] Yeomans J M 1992 *Statistical Mechanics of Phase Transitions* (Oxford: Oxford University Press)
[12] Cavagna A 1999 Irrelevance of memory in the Minority Game *Phys. Rev.* E **59** R3783
[13] Cavagna A, Garrahan J P, Giardina I and Sherrington D 1999 Thermal model for adaptive competition in a market *Phys. Rev. Lett.* **83** 4429
[14] Challet D and Marsili M 2000 Relevance of memory in Minority Games *Phys. Rev.* E **62** 1862
[15] Challet D and Zhang Y C 1998 On the Minority Game: analytical and numerical studies *Physica* A **256** 514
[16] Bottazzi G, Devetag G and Dosi G 1999 Learning and emergent coordination in speculative markets: some properties of 'Minority Game' dynamics *LEM–St. Anna's School of Advanced Studies working paper 1999/24*
[17] Challet D, Marsili M and Zecchina R 2000 Comment on [13] *Phys. Rev. Lett.* **85** 5008
[18] Cavagna A, Garrahan J P, Giardina I and Sherrington D 2000 Reply to [17] *Phys. Rev. Lett.* **85** 5009
[19] Garrahan J P, Moro E and Sherrington D 2000 Continuous time dynamics of the Minority Game *Phys. Rev.* E **62** 9

[20] van Kampen N G 1992 *Stochastic Processes in Physics and Chemistry* (Amsterdam: North-Holland)

[21] Hubbard J 1959 *Phys. Rev. Lett.* **3** 77
Stratonovich R L 1957 *Dokl. Akad. Nauk* **115** 1097

[22] Sherrington D 1967 A new method of expansion in the quantum many-body problem: III. The density field *Proc. Phys. Soc.* **91** 265

[23] Sherrington D 1971 Auxiliary fields and linear response in Lagrangian many-body theory *J. Phys. C: Solid State Phys.* **4** 401

[24] Sherrington D and Kirkpatrick S 1975 Solvable model of a spin glass *Phys. Rev. Lett.* **35** 1972

[25] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8

[26] Challet D, Marsili M and Zecchina R 2000 Statistical mechanics of systems of heterogeneous agents: Minority Games *Phys. Rev. Lett.* **84** 1824

[27] Marsili M, Challet D and Zecchina R 2000 Exact solution of a modified El Farol's bar problem: efficiency and the role of market impact *Physica* A **280** 522

[28] Edwards S F and Anderson P W 1975 Theory of spin glasses *J. Phys. F: Met. Phys* **5** 965

[29] Amit D J, Gutfreund H and Sompolinsky H 1985 Storing infinite numbers of patterns in a spin-glass model of neural networks *Phys. Rev. Lett.* **55** 1530

[30] Sherrington D 1993 Neural networks: the spin glass approach *Mathematical Approaches to Neural Networks* ed J G Taylor (Amsterdam: North-Holland) p 261

[31] Challet D, Marsili M and Zhang Y C 2000 Modeling market mechanism with Minority Game *Physica* A **276** 284

[32] Sherrington D, Garrahan J P and Moro E 2000 Statistical physics of adaptive correlation of agents in a market *Preprint* cond-mat/0010455

[33] Coolen A C and Sherrington D 1993 Dynamics of fully connected attractor neural networks near saturation *Phys. Rev. Lett.* **71** 3886
Coolen A C, Laughton S N and Sherrington D 1996 Dynamical replica theory for disordered systems *Phys. Rev.* B **53** 8184

[34] Martin P C, Siggia E D and Rose H A 1978 Statistical dynamics of classical systems *Phys. Rev.* A **8** 423
De Dominicis C 1978 Dynamics as a substitute for replicas in systems with quenched random impurities *Phys. Rev.* B **18** 4913
Sompolinsky H and Zippelius A 1982 Relaxation dynamics of the Edwards–Anderson model and the mean-field theory of spin-glasses *Phys. Rev.* B **25** 6860

[35] Johnson N F, Hart M and Hui P M 1999 Crowd effects and volatility in a competitive market *Physica* A **269** 1

[36] Hart M, Jeffries P, Johnson N F and Hui P M 2000 Crowd–anticrowd model of the Minority Game *Preprint* cond-mat/0003486

[37] Li Y, VanDeemen A and Savit R 2000 The Minority Game with variable payoffs *Physica* A **284** 461

[38] Farmer J D and Joshi S 1999 Evolution and efficiency in a simple technical trading model *Santa Fe Institute working paper 99-10-071*
Challet D, Chessa A, Marsili M and Zhang Y C 2001 From Minority Games to real markets *Quantitative Finance* **1** 168

[39] Metzler R, Kinzel W and Kanter I 2000 Interacting neural networks *Phys. Rev.* E **62** 2555

[40] Johnson N F, Hui P M, Jonson R and Lo T S 1999 Self-organized segregation within an evolving population *Phys. Rev. Lett.* **82** 3360