

Statistical mechanics of the majority game

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2003 J. Phys. A: Math. Gen. 36 11725

(<http://iopscience.iop.org/0305-4470/36/47/002>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.89

The article was downloaded on 02/06/2010 at 17:16

Please note that [terms and conditions apply](#).

Statistical mechanics of the majority game

P Kozłowski^{1,2} and M Marsili¹

¹ Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11,
34014 Trieste, Italy

² Istituto Nazionale per la Fisica della Materia (INFN), Unità Trieste-SISSA, Via Beirut 2-4,
34014 Trieste, Italy

E-mail: kozl@ictp.trieste.it and marsili@ictp.trieste.it

Received 8 July 2003

Published 12 November 2003

Online at stacks.iop.org/JPhysA/36/11725

Abstract

The majority game, modelling a system of heterogeneous agents trying to behave in a similar way, is introduced and studied using methods of statistical mechanics. The stationary states of the game are given by the (local) minima of a particular Hopfield-like Hamiltonian. On the basis of replica symmetric calculations, we draw the phase diagram, which contains the analogue of a retrieval phase. The number of metastable states is estimated using the annealed approximation. The results are confronted with extensive numerical simulations.

PACS number: 05.20.-y

1. Introduction

The cooperative processes taking place in complex systems of agents interacting with each other, which were up to quite recently studied mainly by human sciences, have also become an interesting object of research for physicists. Despite the tremendous complexity of such systems and the presence of unpredictable ingredients related to human free will, some of the statistical regularities which characterize their collective behaviour can be studied using models and techniques developed in the field of statistical physics.

As usual in such situations one builds very abstract theoretical models which have a much wider applicability. The so-called minority game, for example, was proposed as a formalization of the El-Farol bar problem [1], in order to capture the essential features of traders' interaction in a stock exchange market [2]. But the same model can describe a more general situation where many agents compete for the exploitation of a number of scarce resources [3].

This seemingly simple model turns out to exhibit a surprisingly rich variety of complex behaviour which was studied thoroughly using various methods [4–6]. The statistical

mechanics approach to this sort of models is particularly interesting as it reveals aspects of cooperative phenomena which may be qualitatively different from those studied in physics.

The minority game is based on the assumption that agents prefer to avoid crowds. Hence they tend to use those strategies which let them be as often as possible in the minority. This is only one of the many possible types of interactions. It seems natural to ask what happens under the reversed (majority) rule, favouring those strategies which allow the agents to stay on the side of the majority.

In the economic interpretation of this sort of model, it has been shown [7] that the minority rule describes so-called ‘fundamental’ or ‘contrarian’ traders in a financial market. These are agents who believe that market prices are close to an equilibrium and hence expect that price fluctuations tend to generate price changes towards the equilibrium value. The majority rule is instead appropriate for trend followers, whose behaviour is thought to be responsible for the so-called ‘bubbles’—buy rushes leading to price increases well beyond those which would be justified by an economic evaluation. The majority mechanism is self-reinforcing as it generates self-fulfilling prophecies: if the agents expect that the price will rise (fall), the majority of them will buy (sell) which will actually make the price rise (fall).

Before focusing on the majority game, let us mention that the competition between majority and minority players has been studied in [7] in the simplest setting and in [8] in its full complexity.

From a wider perspective, the majority game describes a situation where the profit of agents increases with the number of agents acting in the same way. Conformity effects of this type are evident in the spreading of fashions. A further example may be that of a shop which lowers the prices as the number of customers increases, or of a product which becomes cheaper the more popular it is. This mechanism, which goes under the name of *increasing returns* in economics, lies at the heart of quite interesting aggregation phenomena—for example, the emergence of cities and economic districts such as Silicon Valley and Hollywood [9].

As we shall see the study of the majority game leads to the analysis of models which are very similar to attractor neural networks, in particular to the Hopfield model [10]. In brief, agents’ learning dynamics is different from Glauber dynamics, but the energy landscape where it takes place is the same. It turns out that aggregation in the majority game is the same phenomenon as memory retrieval in the Hopfield model. Hence the physics of neural networks tells us a lot about the behaviour of the majority game. On the other hand, this study also provides new results on the physics of neural networks by probing energy landscapes such as that of the Hopfield model with a different type of dynamics.

Our work is based on a statistical mechanics approach to the stationary states of the majority game and its results are verified by numerical simulations. The paper is organized as follows: In the next section, we introduce the model and in section 3 we discuss its stationary states—we show that they can be identified with the local minima of the Hopfield-type Hamiltonian. Section 4 deals with the calculation of free energy and the construction of the phase diagram. Given that stationary states are selected in a dynamical way and not according to a Boltzmann weight, we compare our results with extensive numerical simulations. In order to clarify the dynamical behaviour of the model, we discuss the number of stationary states in section 5. We conclude with a summary of the main results and a discussion of their implications.

2. The definition of the model

We consider a system consisting of N agents interacting at discrete time intervals (at each round of the game). The interaction takes place through the action, concerning p objects or

resources, which each agent undertakes. The actions are determined by one of z strategies which are chosen randomly and independently for each of the agents at the beginning of the game. In the course of the game the agents can change their actions only by changing their strategies, i.e. choosing one of the z predefined ones (for each agent). We denote by a_{is}^μ the action taken by agent $i = 1, \dots, N$ concerning resource $\mu = 1, \dots, p$ when he/she adopts strategy $s = 1, \dots, z$. We consider here the case of binary actions $a_{is}^\mu = \pm 1$. The specific values of a_{is}^μ —the realization of quenched disorder—are drawn at random from some distribution. Thus a strategy is a binary vector, which can be interpreted for instance as a list of actions to be undertaken concerning each of the objects.

Agent i chooses the strategy used in the next round of the game on the basis of the performances of his/her strategies in the previous runs, which are measured by score functions $u_{is}(t)$. The agents choose the strategy

$$s_i(t) = \arg \max_s u_{is}(t) \quad (1)$$

with the highest score and undertake the profile of actions $a_{i s_i(t)}^\mu$ $\mu = 1, \dots, p$ ($i = 1, \dots, N$). The agents know about the actions of others only through the cumulative actions:

$$A^\mu(t) = \sum_{i=1}^N a_{i s_i(t)}^\mu \quad \mu = 1, \dots, p \quad (2)$$

which are used to update the score functions:

$$u_{is}(t+1) = u_{is}(t) + \frac{\epsilon}{p} \sum_{\mu=1}^p a_{is}^\mu \left[A^\mu(t) - \eta \left(a_{i s_i(t)}^\mu - a_{is}^\mu \right) \right] \quad (3)$$

where $\epsilon > 0$ and $\eta \in [0, 1]$ are constants.

Let us first discuss this dynamics for $\eta = 0$. Note that $A^\mu(t)$ has the same sign as the action undertaken by the majority concerning object μ . Then equation (3) implies that those strategies prescribing an action aligned with the majority are rewarded. In other words, by this learning dynamics, agents strive to find that strategy which puts them in the majority. Because of the averaging over μ this rule is called a batch version of the majority game. The on-line version, where a value $\mu(t)$ is randomly drawn at each time and agents update the scores depending on $A^{\mu(t)}(t)$, will be discussed in the concluding section.

With parameter $0 < \eta < 1$ one can change the degree to which the agents take into account the influence of their own actions on the cumulative quantity $A^\mu(t)$ (see [11]). In particular, the case $\eta = 1$ describes agents who are learning to respond optimally to the behaviour of the others. Indeed for $\eta = 1$ equation (3) computes the correct value of $A^\mu(t)$ if agent i had actually played strategy s . This is what game theory assumes a rational player should do, so the stationary states of the game for $\eta = 1$ are Nash equilibria (i.e. those states where each agent takes the optimal strategy, given the strategy of others [12]). As in the minority game [4], in spite of the fact that $A^\mu(t) \sim \sqrt{N}$ is much larger than $a_{is}^\mu \sim O(1)$, the η term is not negligible.

The new, updated payoff functions are used to determine the action in the next time step through equation (1).

Note that $A^\mu(t)$ is the difference between the size of the two groups of agents who undertake opposite actions $a_{i s_i}^\mu = +1$ or -1 . If agents do not interact and the actions $+1$ and -1 are equivalent, it is obvious that $A^\mu(t) \sim \sqrt{N}$. We shall pay special attention, in what follows, to the possibility that, when turning on the interaction, a *macroscopic difference* $A^\mu(t) \propto N$ may emerge, for some value of μ .

2.1. $z = 2$ case

In this paper we will focus on the case where $z = 2$. We allow for a correlation of the two strategies of the same agent by introducing a parameter

$$g = P(a_{i+}^\mu = a_{i-}^\mu) \quad (4)$$

with $P(a_{i+}^\mu = 1) = P(a_{i+}^\mu = -1) = \frac{1}{2}$.

Instead of keeping track of the two payoff functions it is enough to consider their difference:

$$y_i = \frac{u_{i+} - u_{i-}}{2}. \quad (5)$$

3. Stationary states

Taking the average over the stationary state distribution, which we denote by $\langle \dots \rangle$, of the dynamical equation of $y_i(t)$, we get

$$v_i \equiv \frac{\langle y_i(t+1) - y_i(t) \rangle}{\epsilon} = \overline{\xi_i} \Omega + \sum_{j=1}^N \overline{\xi_i \xi_j} m_j - \eta \overline{\xi_i^2} m_i \quad (6)$$

where we used a standard notation

$$m_i = \langle \text{sign} y_i \rangle \quad \xi_i^\mu = \frac{a_{i+}^\mu - a_{i-}^\mu}{2} \quad \Omega^\mu = \sum_{j=1}^N \frac{a_{j+}^\mu + a_{j-}^\mu}{2} \quad \bar{x} = \frac{1}{p} \sum_{\mu=1}^p x^\mu.$$

We expect that all the agents in the long time limit are frozen, i.e. they do not change their strategies. Indeed, exactly as in the case of the minority game, it is easy to check that the stationary states correspond to the minima of

$$\begin{aligned} \mathcal{H}_\eta &= -\frac{1}{2} \overline{A^2} + \frac{\eta}{2} \sum_i \overline{\xi_i^2} m_i^2 \\ &= -\frac{1}{2} \sum_{i,j} \overline{\xi_i \xi_j} m_i m_j - \sum_i \overline{\Omega \xi_i} m_i - \frac{1}{2} \overline{(\Omega^\mu)^2} + \frac{\eta}{2} \sum_i \overline{\xi_i^2} m_i^2. \end{aligned} \quad (7)$$

The argument starts by observing that if $v_i \neq 0$, then $y_i \rightarrow \pm\infty$, depending on the sign of v_i , and hence $m_i = \text{sign } v_i$. Only if $v_i = 0$ can we have $m_i \neq \pm 1$. Then one observes that $v_i = -\frac{\partial \mathcal{H}_\eta}{\partial m_i}$ so these conditions are equivalent to the conditions for the minima of \mathcal{H}_η . But it is evident, from the form of \mathcal{H}_η , that its minima lie only at the corners of the hypercube $[-1, 1]^N$.

The conclusion that stationary states correspond to the minima of \mathcal{H}_η is straightforward from the dynamical equations in the limit $\epsilon \rightarrow 0$. Then one can introduce a rescaled continuous time $\tau = \epsilon t$ and verify that \mathcal{H}_η is a Lyapunov function of the continuous time dynamics.

Figure 1 shows that numerical simulations fully confirm the above picture. Note in particular that while the initial stages of the dynamics are somewhat noisy, fluctuations are negligible in the long time limit.

Strictly speaking there is no stationary state in terms of the variables y_i as they diverge to $\pm\infty$. The term *stationary state* refers to the variables m_i which take well-defined values in the limit $t \rightarrow \infty$.

We remark that the \mathcal{H}_0 is simply related to the predictability $H = -2\mathcal{H}_0$ introduced in the minority game [6]. Since $m_i = \pm 1$ for all i , H is also equal to the volatility $\sigma^2 = \overline{A^2}$.

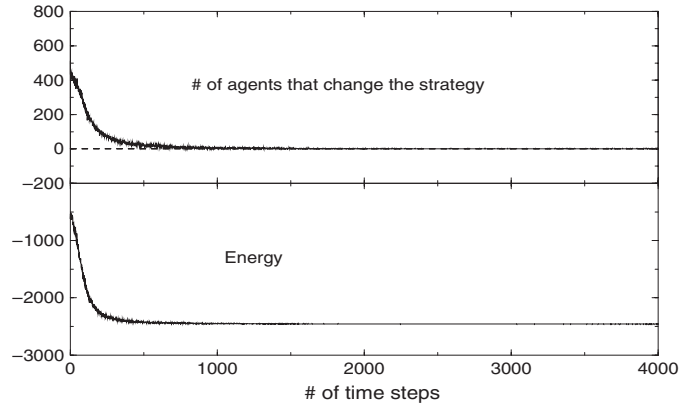


Figure 1. Simulations of the model for $\alpha = 0.2$, $N = 1000$, $\eta = 0$ and random initial conditions.

Hence agents in the majority game strive to maximize H whereas agents in the minority game minimize it.

\mathcal{H}_η is very similar to the Hamiltonian of the Hopfield model known from the theory of neural networks [13]. The only differences are: the scaling with $1/p$ instead of $1/N$, the presence of the random field and the fact that the patterns ξ_i^μ can take three values $0, \pm 1$ instead of only two ± 1 . By analogy with physics we shall call \mathcal{H}_η energy. We note that for $g = 0$ one obtains the ‘pure’ Hopfield model (with different rescaling), whereas for $g = 1/2$ one has a majority game with independently chosen strategies.

The stationary state values of m_i satisfy the equations:

$$m_i = \text{sgn} \left(\frac{\xi_i \Omega}{\xi_i} + \sum_{j=1}^N \overline{\xi_i \xi_j} m_j - \eta \overline{\xi_i^2} m_i \right). \quad (8)$$

It is clear that any configuration $\mathcal{C} = \{m_i\}$ which is a solution of these equations for some value of $\eta \in [0, 1]$ will also be a solution for all $\eta' < \eta$. Hence the set \mathcal{S}_η of stationary states satisfies the property $\mathcal{S}_\eta \subset \mathcal{S}_{\eta'}$ for $\eta' < \eta$ and, in particular, $\mathcal{S}_1 \subset \mathcal{S}_\eta$ for all $\eta < 1$. It is also easy to see that the state with minimal value of \mathcal{H}_η lies in \mathcal{S}_1 for all $\eta \in [0, 1]$. This shows that Nash equilibria are stationary states of the majority game for all values of η , but the converse is not true.

In the remaining part of this paper we will try to determine the nature of the stationary states of the majority game and calculate the number of such states. In the next section we will employ the replica method to analyse the thermodynamics (for $T = 0$) of the model defined by the Hamiltonian (7). Then, using the stability relation (8) we will calculate the number of stationary states in the annealed approximation as a function of various parameters.

4. Replica approach to thermodynamics

In the light of the previous considerations, it is clear that stationary states of the majority game for all values of $\eta \in [0, 1]$ are determined by the minima of \mathcal{H}_η lying at the corners of the hypercube $[-1, 1]^N$. Note that when we restrict ourselves to $m_i = \pm 1$ the η term in \mathcal{H}_η becomes an irrelevant constant. This is why none of the results in this section depend on η .

As usual, we build a partition function corresponding to the Hamiltonian (7), introducing an inverse temperature β . Using the replica method [14], we perform the average over the




Figure 2. The phase diagram computed by the replica symmetric approach described in the text (full line). The points are numerical estimates of the phase boundary $\alpha_c(g)$ obtained from the crossing of the maximal overlap for different system sizes, as shown in the inset ($g = 0.15$, $\eta = 1$, $p = 64, 128$ and 256 , maximal overlap initial conditions).

quenched disorder and obtain the free energy density function. Within the replica symmetric ansatz, the latter depends on the Edward Anderson order parameter q , the overlap b with pattern ξ^1 and the residual overlap r . In the limit $\beta \rightarrow \infty$ we find $q \rightarrow 1$ with

$$\chi = \lim_{\beta \rightarrow \infty} \beta(1 - q)$$

finite and the free energy takes the form

$$f = \frac{1}{2} \frac{b^2}{\alpha} + \alpha \chi r - \frac{\alpha}{2} \frac{1}{\alpha - (1 - g)\chi} - 2\sqrt{\frac{\alpha r}{\pi}} [g + (1 - g) \exp(-b^2/(4\alpha^3 r))] - \frac{(1 - g)b}{\alpha} \operatorname{erf}\left(\frac{b}{2\alpha^{3/2}\sqrt{r}}\right). \quad (9)$$

The parameters b , r and χ should take the values optimizing the free energy. We remark that the overlap b corresponds to the equilibrium value of A^1/N . The capacity $\alpha = p/N$, and g defined by (4) are the parameters.

The analysis of the solutions to the saddle point equations

$$\frac{\partial f}{\partial b} = 0 \quad \frac{\partial f}{\partial r} = 0 \quad \frac{\partial f}{\partial \chi} = 0 \quad (10)$$

allows us to draw the phase diagram (see figure 2).

One can distinguish two phases: a spin glass phase ($b = 0$) and a retrieval phase, where besides a spin glass solution there also exists a retrieval solution ($b \neq 0$).

We also calculated a series of T - α phase diagrams and we found small reentrance of the spin glass phase at low T , just as in the Hopfield model. It implies that as in the Hopfield model there is replica symmetry breaking for low T . However, the effect on the phase diagram is very small and becomes even smaller as g increases.

A simple calculation shows that for the spin glass solution

$$\frac{\overline{A^2}}{N} = \frac{-2\mathcal{H}_0}{N} \simeq \left(1 + \sqrt{\frac{2(1 - g)}{\pi\alpha}}\right)^2$$




Figure 3. Comparison of energy profiles obtained for $g = 0.5$ and $\eta = 0$.

within the replica symmetric ansatz (see figure 3). The solution with $b \neq 0$, which using the neural network nomenclature, we call retrieval, is conveniently described in terms of the parameter $x = b/(2\alpha\sqrt{\alpha r})$ [15], which satisfies the equation

$$x = \frac{(1-g)\text{erf}(x)}{\sqrt{2\alpha(1-g)} + 2\pi^{-1/2}(1-g)[g + (1-g)\exp(-x^2)]}. \quad (11)$$

For $g \leq 2/3$, two non-zero solutions for $x(\alpha)$ exist up to a critical value $\alpha_c(g)$, but only one of them represents a thermodynamically stable state. Of course for $g \rightarrow 0$ we find that the spin glass and retrieval solutions of (10) converge to the solutions found for the Hopfield model [15]. The phase separation line $\alpha_c(g)$ in figure 2 smoothly approaches the $\alpha = 0$ axis $\alpha_c \simeq \frac{75}{2\pi}(\frac{2}{3} - g)^4$ when $g \rightarrow 2/3$.

In the retrieval phase the dynamics can, depending on the initial conditions, end up in one of the two qualitatively different sorts of attractors. In the majority game language the retrieval corresponds to the macroscopic value of $A^\mu \sim O(N)$ for some μ (say 1), whereas $A^\mu \sim \sqrt{N}$ for the remaining values of $\mu = 2, \dots, p$. In the spin glass phase, $A^\mu \sim \sqrt{N}$ for all μ .

To confirm these analytical results we have performed extensive numerical simulations using the dynamical definition of the model (1, 3). This is important because the above calculation is based on the Boltzmann weight and it focuses on the lowest energy minima. There is, however, no guarantee that the dynamics of the majority game selects the minima with the lowest energy.

Direct iteration of the dynamics is slow because in the late stages the time interval between individual spin flips $s_i \rightarrow -s_i$ becomes very large. A much more efficient algorithm is possible in the continuum dynamics for $\epsilon \rightarrow 0$ because then one can integrate easily the dynamics between two consecutive spin flips. Clearly when $t \gg 1$ the continuum time dynamics coincides with the discrete time batch dynamics ($\epsilon = 1$). Indeed no noticeable difference between simulations with $\epsilon = 1$ and $\epsilon \rightarrow 0$ was found in the typical properties of stationary states.

The comparison of the energies obtained using different methods for a few values of α in the spin glass phase is presented in figure 3.

The estimate of the phase separation line coming from the simulations is plotted in figure 2. Each point was obtained from the crossing of the curves A^1/N versus α for different




Figure 4. Overlap A^1/N in the stationary state as a function of initial overlap at $t = 0$ in the retrieval phase ($\alpha = 0.05$ and $g = 0.1$) for $\eta = 0$ (\blacklozenge and \blacktriangle) and 1 (\circ and \square). Two system sizes ($p = 64$ and 128) are shown in order to appreciate finite size effects. Inset: Overlap as a function of α for $\eta = 0$ (\circ , $*$ and \diamond) and $\eta = 1$ ($+$ and \times) for maximal initial overlap. The stable (solid line) and unstable (dashed) solutions of saddle point equations are also shown.




Figure 5. Same as figure 4 but in the spin glass phase ($\alpha = 1$ and $g = 0.1$).

system sizes, as shown in the inset. There is a good agreement with the analytical results. The static results are η independent, but simulations clearly show that this parameter plays an important role in the dynamics.

Figure 4 shows that retrieval states are indeed attractors in the retrieval phase. Even when starting from initial conditions which have only a partial overlap with pattern $\mu = 1$, the dynamics converges to the retrieval state both for $\eta = 0$ and 1 . Actually, retrieval is enhanced when $\eta = 0$.

The inset of figure 4 shows how the overlap A^1/N depends on α , in simulations with maximal initial overlap $A^1(t = 0)/N$. While for large α the overlap essentially vanishes (in the limit $N \rightarrow \infty$) if $\eta = 1$, A^1/N attains a relatively large value for $\eta = 0$. The error bars for $\eta = 0$ indicate that the distribution of the overlap is very broad for intermediate values of α ; the overlap in a particular run can converge to any value in the interval $[0, 1 - g]$.

Figure 5 shows that a very different scenario takes place in the spin glass phase. While for $\eta = 1$ even starting from maximal overlap the dynamics converges to a spin glass state ($\langle A^1 \rangle / N \rightarrow 0$ as $1/\sqrt{N}$ when $N \rightarrow \infty$), for $\eta = 0$ the stationary state preserves the initial

overlap. In order to understand this behaviour it is useful to compute the number of stationary states of the majority game as a function of the parameters α , g and η .

5. Number of stationary states

In order to learn more about the phase space of the majority game we calculate the number of stationary states as a function of various parameters. The stability relation (8) can be written in the form

$$m_i \overline{\xi_i A} - \eta \overline{\xi_i^2} > 0 \quad m_i = \pm 1 \quad (12)$$

The average number of stationary states can be defined as

$$\Phi = \left\langle \left\langle \sum_{\{m_i = \pm 1\}} \prod_{i=1}^N \Theta(m_i \overline{\xi_i A} - \eta \overline{\xi_i^2}) \right\rangle \right\rangle \quad (13)$$

where $\langle \langle \dots \rangle \rangle$ stands for the average over random patterns ξ_i^μ . This quantity is not self-averaging and diverges exponentially with N . Therefore, we will calculate the annealed entropy:

$$s_a = \frac{1}{N} S_a = \frac{1}{N} \ln \Phi. \quad (14)$$

Since the average over the disorder is inside the logarithm (annealed approximation) the calculations are straightforward (see [16, 17]) and lead to the following result:

$$s_a = \max_{c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}} \{s_a(c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma})\} \quad (15)$$

$$\begin{aligned} s_a(c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}) = & c\hat{c} - \alpha\gamma\hat{\gamma} + \alpha^2\Gamma\hat{\Gamma} - \frac{\alpha}{2} \ln[2\Gamma + (\gamma - 1)^2] \\ & + \ln \left[\cosh(\hat{c}) - \frac{1}{2}(1-g) \left[e^{-\hat{c}} \operatorname{erf} \left(\frac{2(1-g)(\eta - \hat{\gamma}) - \frac{2}{\alpha}c}{2\sqrt{2(1-g)\hat{\Gamma}}} \right) \right. \right. \\ & \left. \left. + e^{\hat{c}} \operatorname{erf} \left(\frac{2(1-g)(\eta - \hat{\gamma}) + \frac{2}{\alpha}c}{2\sqrt{2(1-g)\hat{\Gamma}}} \right) \right] - g \cosh(\hat{c}) \operatorname{erf} \left(\frac{2(1-g)(\eta - \hat{\gamma})}{2\sqrt{2(1-g)\hat{\Gamma}}} \right) \right]. \end{aligned} \quad (16)$$

In figure 6 the annealed entropy is plotted as a function of g , η and α . For large values of α the number of stationary states increases dramatically when η decreases from 1 to 0. As α grows s_a saturates at $\ln(2)$ for $\eta \neq 1$. This can be understood observing that when

$$\left| \overline{\xi_i A_{-i}} \right| \equiv \left| \overline{\Omega \xi_i} + \sum_{j \neq i} \overline{\xi_i \xi_j} m_j \right| < (1-\eta) \overline{\xi_i^2} \simeq (1-\eta)(1-g)$$

equation (8) is satisfied for both $m_i = \pm 1$. In words, when the effective field A_{-i} on spin i due to the other spins m_j is weak enough (i.e., its absolute value is smaller than $(1-g)(1-\eta)$), m_i can take both values. If it happens for every i all the states of the system (expressed in terms of configuration $\{m_i\}$) are stationary and thus $s_a = \ln(2)$. For large α the effective fields become small in absolute value. Indeed $\overline{\xi_i A_{-i}}$ is well approximated by a Gaussian variable with zero mean and variance $2|\mathcal{H}_0|/p \sim 1/\alpha$. As a result $s_a \rightarrow \ln(2)$ as $\alpha \rightarrow \infty$. Only when $\eta = 1$ or $g = 1$ does s_a saturate to values smaller than $\ln(2)$ (see figure 7).




Figure 6. α dependence of entropy s_a for various values of the parameters g and η .




Figure 7. α dependence of entropy s_a for $\eta = 1$ and various values of g . Inset: Maximal entropy s_a^{\max} as a function of g for $\eta = 1$ and $\alpha \rightarrow \infty$.

The strong η dependence of s_a displayed in the first plot in figure 6 explains the difference in the dynamical behaviour of the model in the spin glass phase with $\eta = 0$ and $\eta = 1$ (figure 5). Since for $\eta = 0$ the stationary states are very dense in the phase space, the system does not move far from the initial state before it is trapped at one of the fixed points of the dynamics. Thus the initial overlap changes very little. In the case $\eta = 1$ the number of stationary states is much smaller and the system goes far away from the initial state (the initial non-zero overlap vanishes for $N \rightarrow \infty$). It is important to remark that the states with a non-zero overlap which are stable for $\eta = 0$ in the spin glass phase are not attractors in the usual sense because their basin of attraction vanishes in the thermodynamic limit.

It is easy to find s_a as a function of energy E :

$$s_a(E) = \max_{c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}, u} \left\{ s_a(c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}) + uE + \frac{1}{2} \frac{u}{\alpha} c^2 \right\} \quad (17)$$




Figure 8. Entropy $s_a(E)$ as a function of energy for $g = 0.5$, $\eta = 0$ and various values of α .




Figure 9. Entropy $s_a(b)$ as a function of the overlap for $\eta = 0$ and various values of g .

or overlap b

$$s_a(b) = \max_{c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}, u} \{s_a(c, \hat{c}, \Gamma, \hat{\Gamma}, \gamma, \hat{\gamma}) + ub - uc\}. \quad (18)$$

The results are presented in figures 8 and 9.

The energy dependence of s_a enables us to determine the average energy of the infinite system (see the curve in figure 3) and explains the size of the error bars of the simulation results. The larger error bars are due to the wider distribution of the energy (smaller α).

The b dependence of s_a deep in the retrieval phase has a different character than in the spin glass phase (see figure 9). The gap (region where $s_a(b) < 0$) in the distribution in figure 9 disappears close to the phase boundary. We suppose that in reality the gap vanishes precisely at this boundary (compare [17]). Unfortunately, due to the inadequacy of the annealed approximation we are not in a position to draw more quantitative conclusions.

This inadequacy of the annealed approximation used to calculate the entropy s_a is much more evident than in the case of the minority game with $\eta = 1$ [16]. One can see it by comparing the maximal allowed value of the overlap $b_{\max} = 1 - g$ with the maximal value of b for which $s_a(b) > 0$ (figure 9). For all $g > 0$ the function $s_a(b)$ suggests the existence of stationary states with $b > b_{\max}$. Also the discrepancy between the energy profiles (figure 3) should be attributed to the use of the annealed approximation, which indeed underestimates the true energy.

We tried to calculate the quenched entropy using the replica method (see [18]), but up to now we were not able to solve the arising numerical problems.

6. Conclusions and discussion

We have shown that the stationary states of the majority game correspond to the local minima of a Hopfield-type Hamiltonian and are attained when all agents ‘freeze’, i.e. always use the same strategy. Stationary states are not necessarily Nash equilibria except when agents correctly account for their impact on the aggregate ($\eta = 1$). Depending on the parameters, the system can be in one of two phases: a retrieval phase characterized by attractors with a macroscopic overlap $A^1 \sim O(N)$ and a spin glass phase with no retrieval. A macroscopic overlap can also be sustained, in the spin glass phase, for η small. We attribute this phenomenon to the self-reinforcing term $(1 - \eta)\overline{\xi_i^2} s_i$ in the dynamics which causes a dramatic increase in the number of stationary states as η decreases. This picture extends to the on-line version of the game. Indeed the equations for the stationary states are the same and, since agents freeze in the long run, fluctuations play no role (in contrast to the case of the minority game [6]).

These results allow us to draw some suggestions on the behaviour of systems of interacting agents driven by conformity or by increasing returns. The occurrence of a macroscopic overlap $A^1 \sim O(N)$ may correspond to crowd effects such as fashions and trends, when a large fraction of agents behave similarly in some respect, or to economic concentration, when, for example, one particular place is arbitrarily selected for large scale investments. The development of these crowd effects requires: (i) that the number of agents is large compared to the number of resources (α small), (ii) a sufficient differentiation between strategies of agents ($g < 2/3$) and (iii) a large enough initial bias (i.e., an initial macroscopic overlap) towards a particular resource, fashion or place. Finally, crowd effects can be sustained under more general conditions (i.e., in the spin glass phase) if agents do not behave strategically, i.e. if they neglect their impact on the aggregate (η small).

Besides the relevance of the model as a system of heterogeneous interacting agents, it is also interesting as an example of non-Glauber dynamics in the energy landscape of Hopfield-type Hamiltonians. It is remarkable that, in spite of the fact that the dynamics of y_i does not satisfy detailed balance, the statistical mechanics picture remains quite accurate.

Acknowledgments

This work was done within the EU research and training network STIPCO under the contract: HPRN-CT-2002-00319.

References

- [1] Arthur W B 1994 *Am. Econom. Assoc. Papers Proc.* **84** 406
- [2] Challet D and Zhang Y-C 1997 *Physica A* **246** 407
Zhang Y-C 1998 *Europhys. News* **29** 51
- [3] Marsili M, Mulet R and Ricci-Tersenghi F 2003 *Heterogeneous Agents, Interactions and Economic Performance (Lecture Notes in Econ. and Math. Systems vol 521)* ed R Cowan and N Jonard (Berlin: Springer)
- [4] Challet D, Marsili M and Zecchina R 2000 *Phys. Rev. Lett.* **84** 1824
- [5] Coolen A C C and Heimerl J A F 2001 *Phys. Rev. E* **63** 056121
- [6] Marsili M and Challet D 2001 *Phys. Rev. E* **64** 056138
- [7] Marsili M 2001 *Physica A* **299** 93
- [8] De Martino A, Giardina I and Mosetti G 2003 *J. Phys. A: Math. Gen.* **36** 8935
- [9] Krugman P R 1996 *The Self-Organizing Economy* (Oxford: Blackwell)

-
- [10] Hopfield J J 1982 *Proc. Natl Acad. Sci. USA* **79** 2554
 - [11] Marsili M, Challet D and Zecchina R 2000 *Physica A* **280** 522
 - [12] Fudenberg D and Tirole J 1991 *Game Theory* (Cambridge, MA: MIT)
 - [13] Müller B, Reinhardt J and Strickland M T 1990 *Neural Networks; An Introduction* (Berlin: Springer)
 - [14] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
 - [15] Amit D J, Gutfreund H and Sompolinsky H 1987 *Ann. Phys., NY* **173** 30
 - [16] De Martino A and Marsili M 2001 *J. Phys. A: Math. Gen.* **34** 2525
 - [17] Gardner E 1986 *J. Phys. A: Math. Gen.* **19** L1047
 - [18] Bray A J and Moore M A 1980 *J. Phys. C: Solid State Phys.* **13** L469