Dynamical quenching and annealing in self-organization multiagent models

E. Burgos,[1,*] Horacio Ceva,[1,†] and R. P. J. Perazzo[2,3,‡]

[1]*Departamento de Física, Comisión Nacional de Energía Atómica, Avenida del Libertador 8250, 1429 Buenos Aires, Argentina*
[2]*Departamento de Física FCEN, Universidad de Buenos Aires, Ciudad Universitaria–Pabellón 1, 1428 Buenos Aires, Argentina*
[3]*Centro de Estudios Avanzados, Universidad de Buenos Aires, Uriburu 950, 1114 Buenos Aires, Argentina*

We study the dynamics of a generalized minority game (GMG) and of the bar attendance model (BAM) in which a number of agents self-organize to match an attendance that is fixed externally as a control parameter. We compare the usual dynamics used for the minority game with one for the BAM that makes a better use of the available information. We study the asymptotic states reached in both frameworks. We show that states that can be assimilated to either thermodynamic equilibrium or quenched configurations can appear in both models, but with different settings. We discuss the relevance of the parameter G that measures the value of the prize for winning in units of the fine for losing. We also provide an annealing protocol by which the quenched configurations of the GMG can progressively be modified to reach an asymptotic equilibrium state that coincides with the one obtained with the BAM.

I. INTRODUCTION

In recent times considerable attention has been given to the description of self-organization processes in multiagent models. These are geared to explain how (and which) individual decisions give rise to cooperative effects that show up in the behavior of the system as a whole. Such models have a wide number of interesting applications in a variety of situations that range from routing of messages in an information network to the way in which equilibrium is (dynamically) achieved by an ensemble of economic agents.

The main question posed in this line of analysis is how a system composed of multiple agents acting in a decentralized fashion is able to reach a state of ''macroscopic order'' in which individual decisions appear to be coordinated with one another. A well known example of this line of research is the bar attendance model (BAM) proposed by Arthur in Ref. [1]. This model was developed to illustrate a case of coordination through a process that can be considered as inductive learning.

In the BAM a (large) number of individuals have to decide whether or not to go to a bar at a certain date. They share the common perception that attendance is not desirable when the number of people who are present exceeds a critical value, which is smaller than the total population of clients. Otherwise they derive utility from going to the bar.

A dynamics is proposed by which each agent individually updates her criteria to decide on attendance. These are based upon public information, which concerns the total attendance in previous days. Agents use inductive reasoning to predict the number of customers on the following day.

It is simple to realize that the model is self-deceiving: if everybody decides to attend, crowding is produced and everybody loses. Everybody will then refrain from going the next day, a situation in which, again, everybody loses because many could go without crowding. In an iterated version of the model agents plan their attendance for the next few days. In this case it is more evident that the only stable solution is a situation in which a diversity of individual strategies compensate each other, saturating the capacity of the bar every day. The system self-organizes in a (dynamic) Nash equilibrium in which every agent optimizes their respective situation given what has been done by the rest of the players. In this case any player who changes her rules unilaterally ends up losing. Such equilibrium is also Pareto optimal because no agent can improve her situation without harming that of some other player. If adaptation is stochastic, the attendance equilibrates at the accepted value with random fluctuations.

An alternative version of the BAM called the minority game (MG) has recently received much attention [2–6]. The main feature of this framework is that agents derive utility whenever their choice coincides with that of the minority. Self-deception is also evident within this framework. The MG has been presented as a simplified version of the BAM. In the present paper we concentrate on the approach introduced in Ref. [6], making a detailed comparison of the MG and the BAM. Johnson *et al.*'s original model uses the results of several previous rounds of the game (i.e. it has some kind of memory). In earlier work [7], and following the ideas of Cavagna and co-workers [5], we proved that a ''simplified'' version of the model, *that makes no use of memory*, is indistinguishable from the original formulation. In this work we follow the same approach for both the MG and the BAM. We are aware that this point is still controversial in relation to the alternative formulation of Refs. [3,4]. We will see that the MG, besides being a particular case of the BAM (namely, one in which the optimal attendance has been fixed at half the total number of customers), makes a different use of the same information [8].

The treatments of the BAM have traditionally [1,9] been based upon genetic algorithms [10]. Each agent is assumed to keep a set of strings that encode her possible attendance

————
*Email address: burgos@cnea.gov.ar
†Email address: ceva@cnea.gov.ar
‡Email address: perazzo@df.uba.ar

schedules for the following days. After that period each agent computes the outcome of all her possible plans and chooses the best to use during the next period. The attendance schedules are considered as the ''genomes'' of a genetic algorithm and are updated by selection, random mutations, and crossover. The BAM can be regarded as an iterated version of the MG, i.e., one in which players have to make their choices for several rounds of the same game. There is, however, one important difference. The BAM involves only (the iteration of) pure strategies where players either go or do not go to the bar. The players in the MG have instead *mixed* strategies, i.e., they make their choice of going or not going *with a given probability*. The iterated pure strategies in the BAM undergo an adaptive process. Within the MG mixed strategies also change. Each agent randomly attends the bar (with a given probability) and collects points or pays fines depending on whether the attendance is below or above the acceptable level. The mixed strategy (i.e., her attendance probability) is changed when her account falls below a given threshold. Both the BAM and the probabilistic approach of Ref. [6] can be regarded as mean field descriptions because only the aggregate effect of the whole ensemble of agents is relevant for both adaptive processes. On the other hand, the mixed strategy played by the agents in the MG can be assimilated to a time averaged version of the iterated attendance schedules of the players in the BAM.

In order to render possible a comparison between the MG and the BAM two elements have to be considered. On the one hand a generalized version of the MG (the GMG) has to be developed to describe situations traditionally considered within the BAM, namely, cases in which the winning choice is to meet the one taken by an *arbitrary* fraction of the population, not just the minority. There are no systematic studies of this GMG. On the other hand a probabilistic (mixed strategy) mean field approach of the BAM has to be developed. There are a number of possibilities for defining the mixed strategies for these two models. In the present paper we consider for the GMG an obvious extension of the adaptive updating dynamics that reproduces the well known self-organizing features of the MG. For the BAM we propose instead a mixed strategy that reproduces the results obtained with the usual treatments of this model that use iterated, pure strategies. It turns out that the two dynamics are different. It happens that the one associated with the BAM makes a better use of the available information than in the GMG. Although in both frameworks—GMG and BAM—all agents have access to the *same* information, with the proposed BAM rules they are able to keep better track of the causes for winning or losing, thus making a more effective use of it than in the GMG. These seemingly minor changes produce widely different cooperative behaviors. The GMG leads in general to quenched states that are far from equilibrium. The BAM evolution leads instead to an asymptotic configuration that can be thought of as a thermodynamic equilibrium: it has lost memory of the initial conditions, the individual strategies continue to be updated at all times, and yet the density distribution of attendance strategies remains stationary. In the following sections we discuss how both models yield similar results in a rather reduced portion of the space of

parameters but turn out to be very different in all the rest. We show that an essential parameter that governs the self-organized configuration of the system is determined by the rewards and punishments that are contained in the scoring of successes and failures of the agents. The control parameter is the value G of the prize for winning relative to the fine for losing. Ordinarily the effect of this parameter has been disregarded and the only case studied in the literature is $G=1$. Quenching and equilibrium regions change significantly for $G \neq 1$.

In the next section we describe and characterize the quenched state in the GMG and discuss the phase diagram that is found if G is left as a free parameter. In Sec. III we compare this to the BAM. In Sec. IV we provide an analytical discussion to support the results presented in the previous sections. In Sec. V we present an annealing protocol that makes use of the fact that quenching is a consequence of the memory stored in the system encoded by the points gained by the agents. The protocol allows the system to reach equilibrium using the MG dynamics by progressively modifying the quenched state, removing the memory stored in the system.

II. THE GMG MODEL

We consider an N-agent system in which each has a mixed attendance strategy expressed by a probability p_i ($i = 1,2,\ldots,N$) of going to the bar. We take n to be the acceptable level of attendance without crowding. The prescription to adapt the individual strategies is the following. If the attendance A turns out to be $A \leq n$, and the agent chooses to go to the bar, she gains G points. Otherwise, she pays a fine of 1 point. The value G is left as a free parameter.

The balance of gains and losses keeps a record of successes and failures of each agent. If the latter are more frequent than the former, her account ends up by falling below some fixed threshold (which we take to be 0). In this case her account is reset to zero and her strategy p_i is updated. This is done by changing the value of p_i to one chosen at random within the interval $(p_i - \delta p, p_i + \delta p)$. The width of this interval (we take $\delta p = 0.05$) is left unchanged during the whole process. Reflective boundary conditions for p_i are used whenever necessary.

All N agents have initially 0 points and their initial attendance strategies are chosen with some criterion. In the present section we restrict ourselves to considering the case in which all p's are initially chosen at random from the interval $[0,1]$ with uniform probability.

The state of the many agent system is well described by the density distribution $P(p)$ which gives the fraction of agents that have a strategy within the interval $(p,p+dp)$. In order to calculate this we generate an ensemble of H histories of the N-agent system by evolving it during a fixed number of iterations and starting from statistically equivalent initial conditions. The density $P(p)$ is obtained by normalizing the frequency of occurrence of a given strategy in the whole set of the H histories. The number of iterations is fixed by checking that a stationary configuration has been reached.

We also compute other parameters that describe the state

of the system. One is the attendance $A(t)$ averaged over the H histories as a function of the iteration number t; in the following we quote $M(t) \equiv A(t)/N$, to facilitate the comparison for different values of N. A second important parameter is the fraction of agents $C_s(t)$ that change their individual strategies as a function of the iteration number. This fraction provides a direct measure of the stability of the state that has been reached. We will use P_∞ to refer to the asymptotic distribution, $t \to \infty$, whenever the system reaches a stationary state.

A well known result obtained within the framework of the pure MG is a strong polarization of the population. A distribution $P_\infty^{MG}(p)$ is reached in which essentially half of the agents always go to the bar while the other half never go. As a consequence $M(t) \to 1/2$ except for random fluctuations due to the updating of individual strategies that takes place at all times [6,7].

The evolution of a system with the GMG rules and $n \neq N/2$ has been studied by Johnson *et al.* [11] and by us [8]. If the reduced acceptable level of attendance $\mu \equiv n/N$ is close enough to $1/2$ an equilibrium situation is reached that resembles the MG results [see Fig. 1(a)]. Equilibrium is shown by the fact that $M \to \mu$ and $C_s(t)$ stabilizes at a constant, nonzero value after a transient. In Fig. 1(b) we show typical plots of $M(t)$ and $C_s(t)$ for this situation. On the contrary, if μ is larger than a critical value μ_{cr} (see below) a situation is reached in which no further updating of the strategies takes place (*quenched* configuration). In Fig. 2 we show an example of the distribution $P_q(p)$ that is obtained by setting $\mu > \mu_{cr} > 1/2$.

The distribution $P_q(p)$ that is obtained strongly differs from $P_\infty^{GMG}(p)$. The fact that $C_s(t)$ tends exponentially to 0 [see Fig. 2(b)] indicates that a situation has been found that is completely different from the one noted above. This happens because the largest fraction of the population has gained many points and even though their accounts may fluctuate they are not expected to fall below the threshold required to update their individual strategies. The consequence of this is that M stabilizes at a value that may be far from μ. While the situation for $\mu \sim 1/2$ can be thought of as some kind of thermodynamic equilibrium, this case should instead be associated to a quenched configuration in which the state variables are rapidly fixed at values that are close to the initial conditions.

We have performed extensive numerical simulations to investigate this situation using values of μ above and below $1/2$ and several values of G. As expected an overall symmetry is found around $\mu = 1/2$. This is because matching an attendance level greater than $1/2$ is equivalent to matching a symmetric *absence* level equally greater than $1/2$. As will shortly become evident the symmetry is not strict. We discuss this point in Sec. V.

Several features that are displayed in the quenched distribution shown in Fig. 2 are worth noting. In the first place it is found that $P_q(p) = 0$ for $0 < p < 1/2$ and $\mu \gg 1/2$. This can easily be understood. When $\mu \gg 1/2$ the agents that have $p < 1/2$ lose more frequently than they win and therefore update their attendance strategies more frequently, increasing
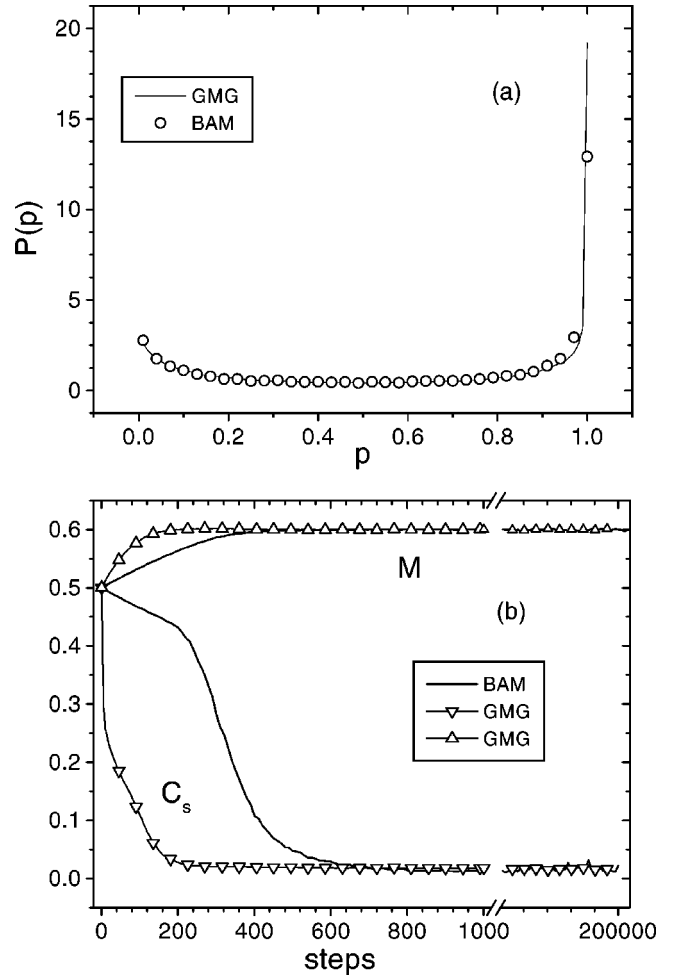

FIG. 1. GMG and BAM without quenching. In this case, $N = 1001$, $G = 1$, and $\mu = 600/1001$. The data are averaged over 200 histories. In (a) we show $P_\infty(p)$ for both cases; (b) illustrates the time evolution of the reduced attendance $M(t)$ and $C_s(t)$, the fraction of moving agents. Notice that $C_s(t)$ does not vanish.

the value of p whenever possible. Once their individual strategies are greater than $1/2$ the inverse happens, agents seldom lose points, and therefore are not forced to update their individual strategies. This also explains why the width of the hump near $p = 1/2$ turns out to be a function of δp only.

The above arguments indicate that the shape of $P_q(p)$ strongly depends upon the initial conditions. This is the reason why we prefer to call this effect ''quenching'' rather than ''freezing.'' The latter suggests a new ordering that bears no relationship to the initial conditions, while the former indicates that a dynamics has been imposed that suddenly reaches a fixed point, leaving the internal parameters of the system at values that are close to the ones chosen with the initial conditions and that may be far from equilibrium.

In Fig. 3(a) we illustrate some cases in which quenching takes place. We plot the asymptotic value of the reduced attendance M reached by the system as a function of μ, displaying only the sector with $\mu > 1/2$, for several values of G.

In Fig. 3(b) we also show a plot of the dispersion σ_M of the values of M obtained in all the H histories of the en-
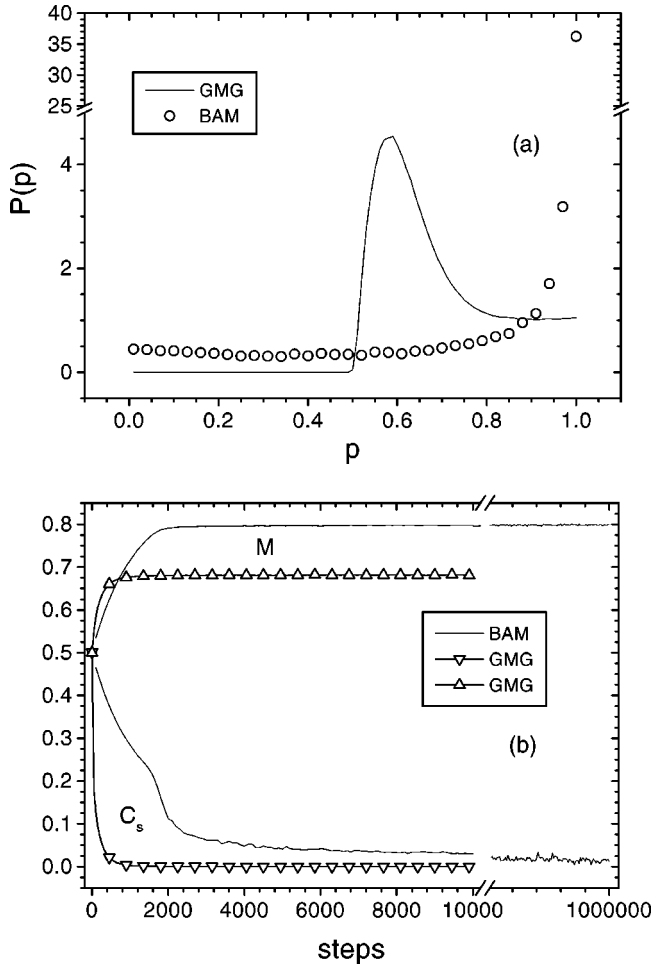
FIG. 2. BAM and GMG with quenching, for the same values of N and G as in Fig. 1, but $\mu = 800/1001$. In (a) we show that the shape of $P(p)$ has little change for the BAM, while the GMG distribution is strongly modified. Correspondingly, in (b) one can see that the GMG does not reach equilibrium ($M \neq \mu$) because the system stops evolving ($C_s = 0$) very rapidly. Data for the GMG are averaged over 3000 histories. On the broken x axis we show that the BAM does not quench even after 10^6 time steps.

FIG. 3. (a) Reduced attendance as a function of μ, for $N = 101$. It can be seen that μ_{cr} is a function of G. Size effects are shown in the inset. (b) Transition from a nonquenched ($\varphi_{20} \neq 1$) to a quenched state ($\varphi_{20} = 1$).

semble $(\sigma_M^2 = \Sigma_h^H [A_\infty(h) - \langle A_\infty \rangle_H]^2 / NH)$. This parameter plays the role of a susceptibility. It displays a peak located at the transition region and is essentially constant outside it. The finite width of the peak as well as the irregularities of M in the transition region are to be attributed to the finiteness of the system.

It is possible to associate a critical value μ_{cr} with the change between the two regimes. In the numerical experiments, the large fluctuations of M near μ_{cr} are due to the fact that some of the histories that make up the ensemble happen to quench while others, purely due to random fluctuations, instead reach equilibrium. As μ increases, the proportions of the two types of history gradually change leading to the situation in which all possible evolutions lead to quenched states. The peaked behavior of σ_M at $\mu = \mu_{cr}$ supports the picture that a sort of critical phenomenon actually takes place in the system for a precise value of μ.

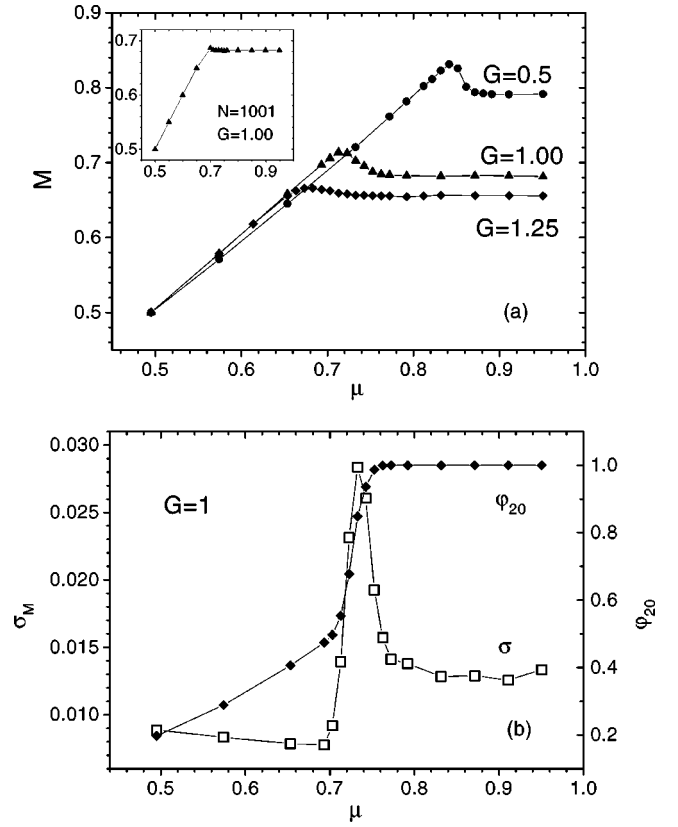There are several criteria to define μ_{cr}. We have chosen

to define it as the minimum value of μ for which the dynamics leads the system to a situation in which the whole population has accumulated a fixed gain, which we (arbitrarily) chose to be not less than 20 points. In Fig. 3(b) we show φ_{20}, the fraction of the population that has accumulated at least 20 points, as a function of μ. It is clear that if this happens any further changes in the distribution $P(p)$ or in the attendance are extremely unlikely to occur. Another possibility is to choose the value of μ that corresponds to the intersection of the lines that are associated with the two regimes or the maximum of σ_M. Again, the differences between these definitions are to be attributed to the finiteness of the systems that we have studied.

The value of μ_{cr} depends upon both G and N. Changes in G have not been explored previously and cause important changes in the region in which quenching occurs as well as in the shape of the distribution $P_\infty^{GMG}(p)$. The effects of the value of G on μ_{cr} can be seen in Fig. 3(a). The general phase diagram (see Fig. 4) shows two regions that correspond to quenched and unquenched systems. The border between the regions is given by the function $\mu_{cr}(G)$. Finite-size effects are also displayed in this figure through the values of $\mu_{cr}(G)$ for two values of N. The results displayed in Fig. 4 allows us to conjecture that for an infinite system with $G = 1$ and $\mu = 1/2$ the usual settings of the minority game lead to an unstable quenched phase.
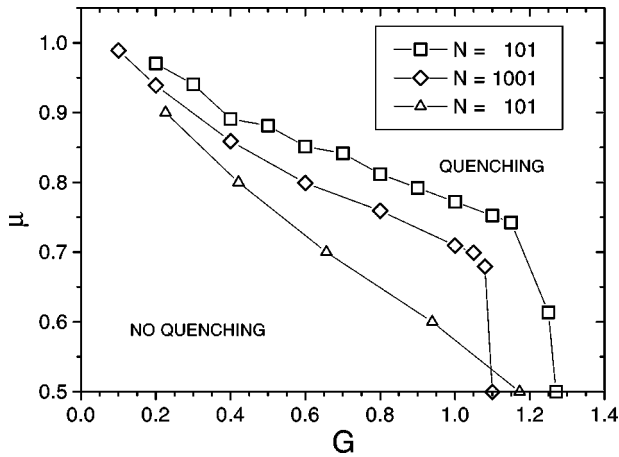
FIG. 4. Phase diagram for the GMG in the μ-G plane. $\triangle$ shows an analytical solution for a δ-type initial condition, while $\square$ and $\diamond$ correspond to numerical solutions starting from a uniform distribution. The critical border for the BAM with δ-type initial conditions would be a vertical line at $G \simeq 1,22$ for $N = 100$.

The changes in the distribution $P_\infty^{GMG}(p)$ that are due to changes in μ keeping G fixed are shown in Figs. 1, 2, and 5. An increase in μ produces an increasingly asymmetric distribution. This hold up to the point at which $\mu > \mu_{crit}$ and quenching occurs, giving rise to a humped distribution $P_q^{GMG}(p)$. The changes introduced by different values of G are shown in Fig. 5. The overall effect of this parameter is to increase the polarization, preserving the level of asymmetry fixed by μ. Values of $G < 1$ are reflected in distributions that are barely peaked at the corners. The polarization increases for $G = 1$ and values of G that are only slightly larger than 1 produce distributions $P_\infty^{GMG}(p)$ that vanish for all values of p except in the close neighborhood of $p = 0$ and $p = 1$. Values of G that are even larger place the system into the quenching region of the phase diagram because every successful initial move produces such a large gain that a revision of the strategy becomes extremely unlikely; therefore the evolution of the system produces only mild changes in the initial distribution.
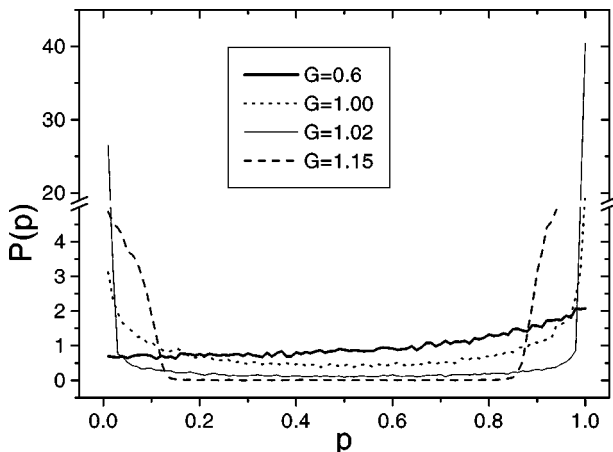

FIG. 5. Change of $P(p)$ with G, for $\mu = 0.6$.

III. THE BAM

The dynamics of the GMG does not keep track of the situations that lead to successes or failures. However, this is stressed in the BAM. It has been customary to accept that agents derive utility in two situations: one when they decide to go to the bar and $M \leqslant \mu$ and also when they refrain from going to the bar and $M > \mu$.

We propose for the BAM an adaptive dynamics that is based upon the same information as in the GMG but uses it differently in order to keep track of the situations that lead to gains or losses, i.e., by correlating them with the action of going or not going to the bar. To this end we introduce ''points'' and ''credits.'' Points are gained or lost if the option of *attending* the bar is correct or not. Credits are instead gained or lost depending whether the option of *not attending* the bar is correct or not. We therefore carry a double account. If an agent chooses to attend (not to attend) the bar and $M < \mu$ $(M > \mu)$ she gains a point (a credit). Otherwise, she loses a point (a credit). When either the number of points or the number of credits falls below some threshold, the corresponding account of points or credits is reset to zero and the strategy is changed in the same fashion as already mentioned for the GMG.

Note that the attendance strategy of an agent is changed by either of two reasons. Such an updating rule makes use of information that is lost with the scoring rules used for the GMG. Suppose, for instance, that an agent having initially all her accounts at zero chooses to go to the bar one day on which $M < \mu$ and chooses not to go the next day when the attendance turns out to be again $M < \mu$. Within the GMG rules she ends up with zero points and her strategy is not updated. With the double account of the BAM she ends up instead with one positive point and one negative credit and is therefore forced to update her strategy [12].

We have performed numerical simulations using this dynamics with similar settings as those reported in the preceding section. The results for the BAM are compared with those obtained with the GMG dynamics in Figs. 1 and 2. When quenching is absent in the GMG, the asymptotic distributions $P_\infty^{BAM}(p)$ can barely be distinguished from $P_\infty^{GMG}(p)$. For this case the values of the attendance $A(t)$ and $C_s(t)$ for the BAM are shown in Fig. 1(b). A strong difference occurs, however, when $\mu > \mu_{cr}$ and the GMG gets trapped in a quenched configuration. For the same settings the BAM converges to equilibrium instead. The asymptotic distribution P_∞^{BAM} shows a strong polarization and the attendance reaches the accepted level without difficulty. The nature of the asymptotic state has no difference from that for $\mu < \mu_{cr}$. This can be seen in Fig. 2(b).

A more profound difference between the two models arises when the above analysis is repeated for $G \neq 1$. While the GMG dynamics can be checked to lead to quenched configurations for essentially all values of G, as shown in Fig. 4, the BAM displays no quenching independently of the value of G. Quenching can, however, occur in the BAM but for different initial conditions from a uniform distribution. We turn to this point in the next section.

IV. ANALYTICAL DISCUSSION

In this section we discuss analytically the relevant features of the density distribution $P(p)$ as well as the occurrence of quenching within both the GMG and the BAM.

A. The density distribution

The shape of the distribution $P(p)$ is the combined result of the actions of the whole ensemble of agents. Each player adjusts her attendance, attempting to minimize her individual loss. When equilibrium is reached the resulting $P_\infty(p)$ for both GMG and BAM concentrates the population in the immediate neighborhood of $p \simeq 0$ and $p \simeq 1$, plus an almost vanishing contribution from intermediate values. The ratio of the areas below the two peaks is, essentially, equal to $\mu/(1-\mu)$. An intuitive guess for $P(p)$ could well be instead a sharply peaked function centered at $p \simeq \mu$. The difference has to be sought in the properties of the distribution $\mathcal{P}(A)$ that gives the probability of occurrence of a party of A customers attending the bar. The distribution $\mathcal{P}(A)$ is of course completely determined by $P(p)$. If $P(p) \simeq \delta(p-\mu)$ the corresponding $\mathcal{P}(A)$ turns out to be large for values of A that differ from the optimal one; the highly polarized equilibrium distribution that is actually generated produces instead a $\mathcal{P}(A)$ that is much more concentrated near $A = N\mu$.

Let us assume without loss of generality that all the agents distribute themselves into D different strategies p_i, $i = 1,\ldots,D$ according to

$$P(p) = \sum_i^D \frac{n_i}{N} \delta(p-p_i), \tag{1}$$

where $\Sigma_i^D n_i = N$. With this assumption, the distribution $\mathcal{P}(A)$ can be written as

$$\mathcal{P}(A) = \sum_{\ell_1=0}^{n_1} \cdots \sum_{\ell_D=0}^{n_D} \prod_i^D \left[\binom{n_i}{\ell_i} p_i^{\ell_i} (1-p_i)^{n_i-\ell_i} \right]$$
$$\times \delta\left(A - \sum_j^D \ell_j \right). \tag{2}$$

The value $(A-N\mu)^2$ measures the departure from the optimal attendance, and therefore $\mathcal{C}^2 = \Sigma_A \mathcal{P}(A)(A-N\mu)^2$ can be taken to measure the square of the average cost incurred by the ensemble of agents while adjusting their individual strategies [13].

Although Eq. (2) is in general difficult to evaluate, $\mathcal{C}^2$ can be calculated without great difficulties. In fact if its definition is used together with Eq. (2), the summation over A can be performed first, taking advantage of $\delta(A-\Sigma_j^D \ell_j)$, and all the summations over the ℓ_j indices decouple from each other. We get

$$\mathcal{C}^2 = \sum_{\ell_1=0}^{n_1} \cdots \sum_{\ell_D=0}^{n_D} \prod_i^D \left[\binom{n_i}{\ell_i} p_i^{\ell_i} (1-p_i)^{n_i-\ell_i} \right]$$
$$\times \left(N\mu - \sum_j^D \ell_j \right)^2. \tag{3}$$

Moreover, from Eq. (1), $\langle p \rangle = \Sigma_i n_i p_i/N$ and $\langle p^2 \rangle = \Sigma_i n_i p_i^2/N$; thus in the limit $D \to \infty$ one gets

$$\mathcal{C}^2 = N^2(\mu - \langle p \rangle)^2 + N(\langle p \rangle - \langle p^2 \rangle). \tag{4}$$

In Eq. (4) $\langle p \rangle$ stands for $\int p P(p) dp$. If $P(p) = 1/N$, as for the initial conditions chosen in all the simulations in the preceding sections, the dispersion turns out to be $\mathcal{C}^2 = N^2(\mu - 1/2)^2 + N/6$. We thus see that such an initial condition is a good guess for the traditional settings of the MG when $\mu \simeq 1/2$, but is indeed very poor for the GMG when $\mu \neq 1/2$.

The equilibrium result $\langle p \rangle = \mu$ that is always found is seen to eliminate the $O(N^2)$ terms in $\mathcal{C}^2$. This is also achieved with the guess $P(p) = \delta(p-\mu)$. However, all $O(N)$ terms do not cancel because $\mathcal{C}^2 = N\mu(1-\mu)$. If the two-peaked equilibrium distribution is roughly approximated by $P(p) = (n_1/N)\delta(p-p_1) + (n_2/N)\delta(p-p_2)$ one readily sees that the $O(N^2)$ terms are eliminated when $n_1 p_1 + n_2 p_2 = \mu N$. The $O(N)$ terms also cancel if the two peaks are $p_1 = 0$, $n_1 = N(1-\mu)$ and $p_2 = 1$, $n_2 = N\mu$.

From this point of view, quenched states are far from optimizing the average aggregate cost $\mathcal{C}^2$ because the corresponding $P_q(p)$ do not even cancel the $O(N^2)$ terms. The relaxation dynamics that tends to minimize individual losses is thus seen to also optimize a global parameter ($\mathcal{C}^2$) only if equilibrium is reached. Otherwise the relaxation stops when no agent loses. This situation, in general, does not corresponds to a minimal average cost.

B. BAM

In this section we discuss the occurrence of equilibrium and quenching within the BAM. This can only be done without any approximation if one assumes a distribution $P(p) = \delta(p-p_o)$. This is seen to be appropriate except for $O(N)$ terms in the dispersion of party sizes.

Given an acceptable level of attendance n, if all agents have the same attendance strategy p_o, the probability of occurrence of a party of *less than* n agents is $S(N,n,p_o)$ where

$$S(N,n,p_o) = \sum_{i=0}^{n-1} \binom{N-1}{i} p_o^i (1-p_o)^{N-1-i}. \tag{5}$$

This is a continuous and monotonically decreasing function of p_o and satisfies the following three conditions:

$$S(N,n,0) = 1 \quad \forall n, \tag{6}$$

$$S(N,n,1) = 0 \quad \forall n, \tag{7}$$

$$S(N,n+m,p_o) \geq S(N,n,p_o) \quad \forall p_o, m > 0. \tag{8}$$

If an agent decides to go to the bar, her account in points will remain equilibrated when the expected reward equals the expected fine, i.e., when the attendance strategy p_p that is common to all the agents satisfies:
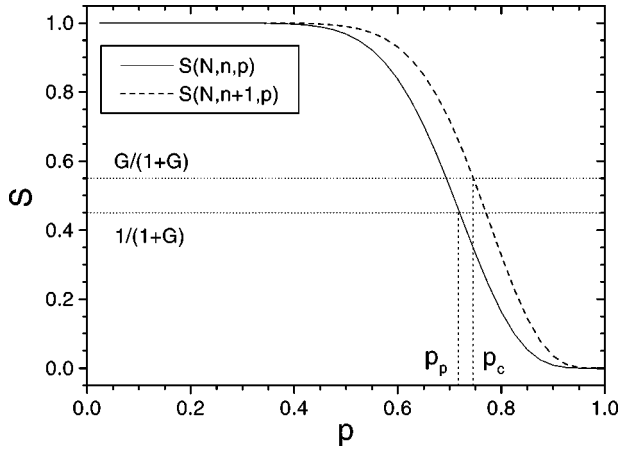
$$S(N,n,p_p) = \frac{1}{1+G}. \tag{9}$$

FIG. 6. The roots p_p and p_c of Eqs. (9) and (10) are exemplified for $N=20$ and $n=14$. These values are chosen for the sake of clarity of the figures. Notice that as G increases, p_c becomes smaller and p_p becomes larger.

If an agent decides instead not to go to the bar a similar condition can also be written for her account in credits. This is

$$S(N,n+1,p_c) = \frac{G}{1+G}. \quad (10)$$

Given G, N, and n, the two equations (9) and (10) define the two roots p_p and p_c that are, respectively, the probabilities with which all agents have to go if they hope to equilibrate, respectively, their accounts of points and credits.

Assume now that $G=1$. Bearing in mind Eq. (8) the two roots satisfy $p_p < p_c$. If the agents choose to attend the bar with a probability $p_o < p_p$, they gain points. On the other hand, if they choose to attend with probability $p_o > p_c$ they gain credits. However, due to the fact that $p_p < p_c$ if they gain points they lose credits and vice versa. In either case all agents will continue to update their strategies, quenching is not possible, and equilibrium is eventually reached.

For a δ-type distribution of strategies the only chance to have a quenched phase using the BAM dynamics is to find a situation in which $p_c \lesssim p_p$. For finite N this can only happen if $G>1$ (see Fig. 6). From the above two equations it is easy to see that the root p_c becomes smaller for increasing values of G and the opposite happens for p_p. A value G_{crit}^{BAM} can eventually be found at which both roots satisfy $p_p = p_c \equiv p_{crit}$. For any $G > G_{crit}^{BAM}$ the regions for winning points and credits start to overlap and quenching occurs within a narrow overlapping interval (p_c, p_p). In Fig. 7 we show the two roots p_p and p_c as functions of G. It further turns out that for finite values of N the value of G_{crit}^{BAM} is nearly independent of μ.

The evolution of a δ-type distribution centered at a value of p_o depends upon the choice of G and μ. If $G > G_{crit}^{BAM}$ and p_o is inside the quenching interval (p_c, p_p) the distribution will not change and the attendance M will remain fixed at the value chosen for the initial conditions. This is so because all agents win both points and credits. If p_o lies outside the
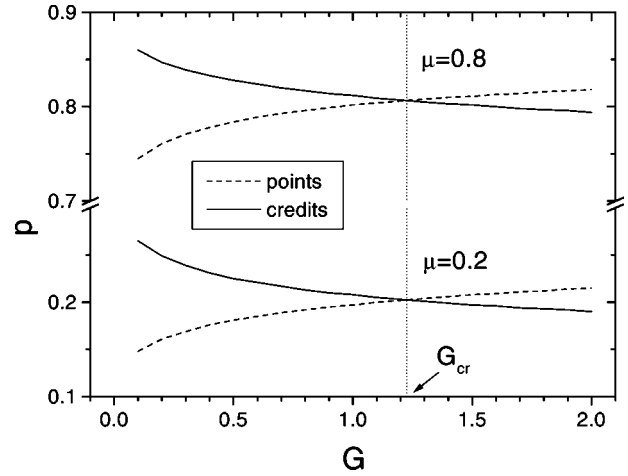

FIG. 7. Quenching for the BAM for δ-type initial conditions is possible for $G > G_{cr}$. The roots p_p and p_c are plotted as functions of G and the indicated values of μ. The two roots are equal for $G = G_{cr}$. The value of G_{cr} is largely independent of μ.

quenching region, agents are forced to update their strategies either because they lose points or because they lose credits. The average attendance M will approach the value μ. Note that the average attendance and the average of individual strategies are completely equivalent, i.e., $M = \int pP(p)dp$. Therefore when the average attendance approaches the quenching interval (p_c, p_p) the evolution will progressively come to a stop. For values of G up to $G \sim 5G_{crit}$ the quenching interval (p_c, p_p) is always small ($|p_p - p_c| \sim 0.1$) and the convergence to a quenched phase is in general very slow. As the mean of the distribution approaches the quenching interval, a smaller fraction of agents have to update their strategies and therefore the fraction $C_s(t)$ tends slowly to zero, very nearly approaching a power law, for a considerable number ($\sim 10^6$) of iterations. Since G_{crit}^{BAM} is almost independent of μ this behavior does not change with its particular value. In the cases in which $p_o \notin (p_c, p_p)$ the corresponding (asymmetric) distributions $P_q^{BAM}(p)$ associated with quenched configurations look similar to those associated with equilibrium because they involve agents with all possible attendance strategies.

C. GMG

A similar analysis as the one performed for the BAM can be made for the GMG. If the system is prepared with a distribution $P_0(p) = \delta(p - p_o)$ the single account of points remains equilibrated provided that the probability of winning $P_w(N,n,p)$ satisfies:

$$P_w(N,n,p_o) = p_o S(N,n,p_o) + (1-p_o)[1 - S(N,n+1,p_o)]$$

$$= \frac{1}{1+G}. \quad (11)$$

In Fig. 6 we show an example of the functions $S(N,n,p)$, $S(N,n+1,p)$, and $P_w(N,n,p)$ for typical values of N and n, and $n > N/2$.

The situation for the GMG is completely different from the BAM. The function $P_w(N,n,p)$ has one maximum in the interval $0 \leq p \leq 1$, and to a high degree of accuracy displays the symmetry $P_w(N,n,p) = P_w(N,N-n,1-p)$. The symmetry, however, is not strict. To see this note that if $n = N$ going to the bar is always a winning option no matter the option of the rest of the agents. If $n = 0$, the action of not going is not enough to win, because that depends upon the choice made by all the other agents.

A critical value G_{cr}^{GMG} and a critical probability of attendance p_{crit}^{GMG} can be defined for the GMG as the p value that corresponds to the maximum of $P_w(N,n,p)$ given by Eq. (11) through

$$P_w(p_{crit}^{GMG}) = P_w^{max} = \max_p P_w(N,n,p), \qquad (12)$$

$$G_{cr}^{GMG} = \frac{1}{P_w^{max}} - 1. \qquad (13)$$

For given values of n and N quenching is produced if $p = p_{cr}$ and $G \geq G_{cr}$. The values of G_{cr} for the GMG are completely different from those of the BAM. While in the latter there exists a common value of G_{cr} independent of μ this is not the case for the GMG (see Fig. 4).

Above G_{cr} and with the initial condition $P_0(p) = \delta(p - p_1)$, the quenching condition is satisfied for any p_1 in the interval $p_< < p < p_>$, where $p_<$ and $p_>$ are the two roots of Eq. (11). Any system prepared with p_1 within that region will rapidly be quenched. Assume for the sake of concreteness that we choose $\mu > 0.5$. The system evolves quite differently if $p_o < p_<$ or $p_o > p_>$. In the first case the whole population of agents is forced to update their strategies to change the individual p's in order to surpass the value $p_<$. Once this happens, updating progressively comes to a stop as the average attendance gets into the interval $(p_<,p_>)$ and the system gets quenched with a value of M below the accepted fraction μ. Since $P_w(N,n,p)$ is such that $p_>$ is very nearly similar to μ, for $p_1 > p_>$ the initial condition of a δ-type distribution changes to $P_\infty^{GMG}(p)$ with a mean value that is very close to μ; this cannot be distinguished from an equilibrium distribution. In addition, updating of strategies continues to take place continuously. Examples are shown in Fig. 8.

D. Finite-size effects

An estimate of the finite-size effects of the models on the values of G_{cr} and p_{cr} can easily be made for the usual settings of the MG, by assuming a δ-type distribution as an initial condition. For $\mu = 1/2$ and from Eqs. (11) and (13) it follows immediately that $p_{crit} = 1/2 \; \forall N$ and the value of G_{crit}^{GMG} can directly be estimated from the value of $P_w^{GMG}(p = 1/2)$ using Eq. (11) and standard asymptotic expansions. It turns out that

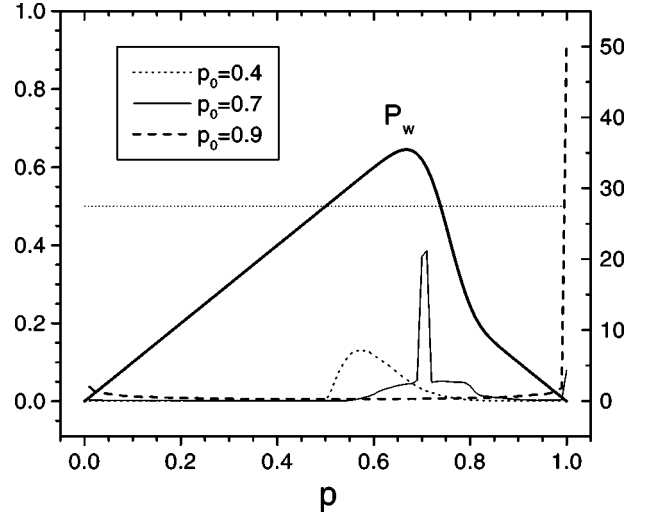$$G_{cr}^{GMG} = \frac{\sqrt{[2\pi N]} + 2}{\sqrt{[2\pi N]} - 2}. \qquad (14)$$

FIG. 8. Distributions $P(p)$ for the GMG (right axis), obtained from three initial conditions $\delta(p - p_o)$. We also show P_w (left axis) to help to understand the evolution of the systems. The roots $p_<$ and $p_>$ correspond to the two roots $P_w = 1/2$. In this case, $\mu = 75/101$.

If this result remains valid for a uniform initial distribution, the usual settings for the MG ($G = 1$ and $\mu = 1/2$) correspond to a critically quenched system.

In Fig. 9(a) are shown $G_{crit}^{GMG}(\mu = 1/2)$ and $G_{crit}^{BAM}(\mu = 1/2)$ as a function of $1/N$. It is seen that $G_{crit}^{GMG} \simeq G_{crit}^{BAM}$ as $N \to \infty$. The corresponding values for $\mu \neq 1/2$ can be related to those shown in Fig. 9(a). In the first place one should notice that the value of G_{crit}^{BAM} is independent of μ. On the other hand, the curves of $G_{crit}^{GMG}(\mu \neq 1/2)$ can be reduced to the one shown through $G_{crit}^{GMG}(\mu) = (a\mu + b)G_{crit}^{GMG}(\mu = 1/2)$ to a very good degree of approximation [a and $b \sim$ constants, up to terms $O(1/N)$].

Concerning the values of p_{crit} one can check [Fig. 9(b)] that p_{crit}^{BAM} and p_{crit}^{GMG} approach each other and tend to the common limiting value $p_{crit}^\infty = \mu$, as expected from a naive estimate of the probability of attendance. One can also conjecture that $G_{crit}^{BAM}(\infty) = 1$.

V. ANNEALING

The situations described hitherto indicate that outside a very restricted region of μ and for different initial conditions the dynamics of the GMG rapidly approaches quenching. The many agent system stops evolving, the attendance remains fixed at a value that is far from the tolerated fraction μ, and all the other internal parameters remain close to the initial conditions. In contrast, the BAM always approaches a distribution $P_\infty^{BAM}(p)$ that corresponds to what one expects of thermodynamic equilibrium: the attendance fluctuates around the expected fraction μ, agents continue to update their individual strategies, and $P_\infty^{BAM}(p)$ remains stationary.

There is a way to regain equilibrium within the framework of the GMG through an iterative procedure that can be regarded as an annealing protocol. The key point is to realize
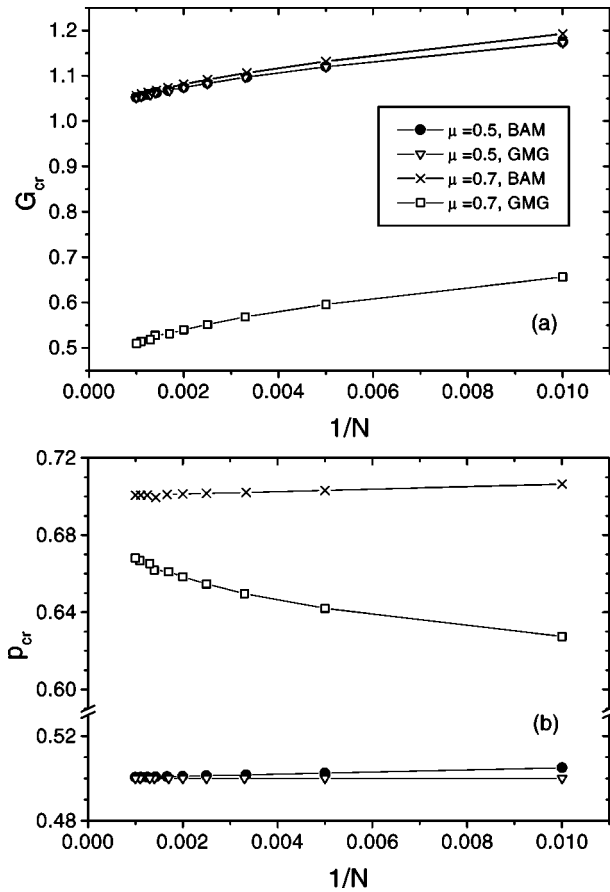
FIG. 9. (a) G_{crit} vs $1/N$ up to $N=1000$. Equation (14) provides an excellent fit that cannot be distinguished from numerical data for the GMG, $G=1, \mu=0.5$; (b) p_{cr} as a function of $1/N$ for the same values of μ.

that quenching is produced by the memory stored in the points gained by the agents. In fact the quenched state remains unchanged because on the average agents continue to gain points. To correct the quenched phase, a procedure can be implemented that consists in periodically obliterating the memory of the system. This is done by removing all the points gained by all the agents in each iteration. In each annealing episode the system is allowed to evolve freely during a maximum of n_s steps; at the end, several parameters are calculated: the fraction φ_5 of agents that have gained not less that 5 points, the relative attendance M, the fraction of histories f_h where the system does not reach a quenched configuration in the allowed n_s steps, and the fraction $f_q = n_q/n_s$, where n_q is the actual number of evolution steps that are needed to quench the system.

The results obtained are shown in Fig. 10 for two values of G and $\mu=850/1001$. At the end of the first episode the distribution P_q [Fig. 10(a)] that is found is entirely similar to the humped distribution shown in Fig. 2, and $\varphi_5=1$, thus corresponding to a quenched phase. As agents are repeatedly deprived of points some of them are increasingly forced to update their strategies. For $G=1$ [Fig. 10(b)] and during the first 30 annealing episodes very few steps are required in order to have $\varphi_5=1$; after ~35 annealing episodes a clear
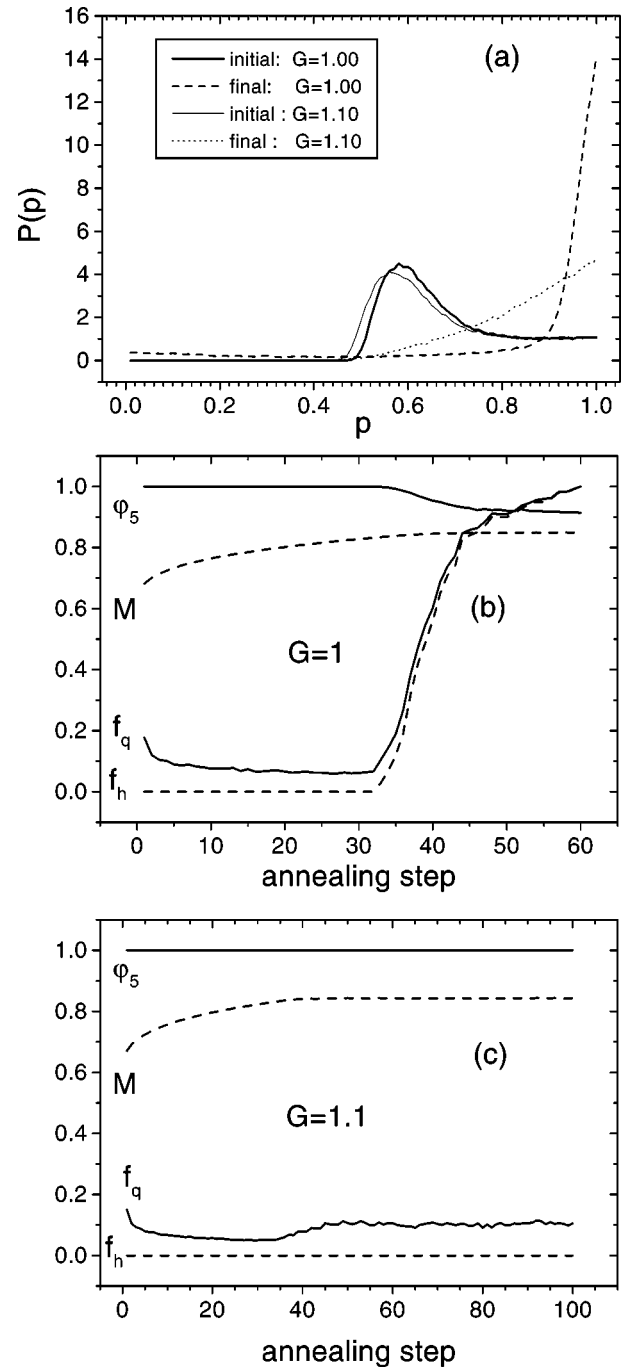

FIG. 10. Annealing in the GMG.

transition takes place that can be interpreted as a gradual melting of the system: the fraction φ_5 drops while f_h and f_q grow abruptly, almost reaching their maximum possible values. For $G=1.1$ [Fig.10(c)] this transition is not present as can be seen from the fact that φ_5 is constantly equal to 1, while $f_h=0$. It is most remarkable, however, that the relative attendance M still reaches the value μ, as for $G=1$. In other words, the annealing protocol is seen to be able to drive the system toward $M \simeq \mu$ with or without a change of the phase of the system.

VI. CONCLUSIONS

In the present work we have discussed the self-organization properties of two multiagent models, the BAM and the GMG. We have shown the relevance of two control parameters. One is the fraction of the population of agents μ that has to combine in order to make the correct decision. The other is G which represents the gain produced by a correct decision, measured in units of the fine paid for making the wrong one. For the sake of concreteness we have borrowed the framework of the BAM in which the decision that has to be made is to go or not go to a bar. The decision is the correct one if the fraction of the population that has decided to go does not exceed μ.

The system under consideration lacks any interaction or internal correlation among the agents. Each agent feels the presence of the rest only through the value of an aggregate parameter that is the total attendance, and rewards or penalties of individual actions are decided with respect to its particular value at each moment. Otherwise, agents act independently from each other. With these rules, the many agent system adapts and always reaches stable asymptotic configurations. These are, however, of quite different nature depending upon the values of the control parameters and of the initial conditions, and on the information that is used to guide the individual actions.

The usual framework of the minority game, if generalized to take into account arbitrary values of μ, gives rise either to states of equilibrium or to quenched configurations. The former have a density distribution of individual strategies that depends only upon the values of G and μ and remains stationary while a fraction of the population C_s continuously updates their strategies. Quenched states have instead frequency distributions that may be very different depending upon the initial conditions, while C_s drops exponentially to zero in all cases.

A phase diagram for the GMG can be established in the μ-G plane. Quenching occurs outside a region limited by a critical line $\mu_{cr}(G)$. This borderline has been obtained numerically for uniform initial conditions. If a δ-type distribution is assumed analytical expressions can be obtained. In both cases finite-size effects are found. A rigorous estimate can be made for an infinite system in the second case. It is found that if $N \rightarrow \infty$ the point ($G=1, \mu=1/2$) belongs to the critical line. A numerical estimate of finite-size effects for a uniform initial distribution is also in agreement with this result. It is interesting to note that these are the usual settings for the minority game widely used in the literature.

The shape of the distribution $P(p)$ is the combined result of the actions of many ''independent'' agents. In spite of such independence we have proved that the equilibrium distribution optimizes the statistical occurrence of the most convenient coordinated state or, equivalently, minimizes the average aggregate cost of the ensemble of agents.

A quenched state is reached when the individual updating of strategies stops. This should not be considered to correspond to a ''frozen'' configuration, however. If this were the case, an internal ordered state would be found bearing no relationship to the initial conditions. Contrary to this, the quenched state is far from being unique; the ''glassy'' structure of the quenched phase lacks any preferred internal ordering. This should be regarded as the signature of the absence of interactions among the agents and has to be attributed purely to dynamic effects. While equilibrium is associated with the optimization of the global magnitude, this is no longer true for a quenched configuration.

Reaching a quenched phase is a direct consequence of the evolution dynamics contained in the rules for updating the individual strategies. This effect is evident when the GMG is compared to the dynamics arising from the BAM with mixed strategies, as we have presented here. In contrast to the GMG, when starting from a uniform distribution the BAM always reaches a state of equilibrium, independently of the values of μ and G. The two models differ in the use of the available information. While the GMG dynamics only keeps a record of good and wrong decisions, the seemingly minor modifications implied in the BAM rules allow it to keep track also of the *reasons* for such successes and failures. The system thus keeps a better ''memory'' of the previous steps in the dynamics. The consequence of this is that the system always self-organizes in such a way that $M \rightarrow \mu$.

The many agent system that has been presented here has several features that are reminiscent of the physics of many body systems. In the first place one should note that a quenched phase can be induced for critical values of μ and G. A sort of order parameter can be established with the relative attendance $M(t)$. Two regions can be individualized in parameter space, one in which $M(t) \rightarrow \mu$ and the other in which $M(t) \rightarrow M_\infty \neq \mu$. For critical values of μ and G, the system undergoes an internal transformation reflected by a peak in the dispersion σ_A that recalls the behavior of the susceptibility at a second order phase transition.

Although the model does not contain any two-body interaction, within this analogy ''particles'' accommodate themselves following an occupational constraint (the maximum allowed attendance) very much in the same fashion as in a many fermion system. To support this viewpoint, one could note that very few ''particles'' participate in the dynamics. The parameter μ fixes the total attendance and therefore plays the role of the chemical potential through an effective collective interaction. From this point of view individual gains should be related to (the negative of) a kind of single particle energy: those agents that have accumulated large gains are less likely to participate in the dynamics.

A resemblance can also be established between quenching in the GMG and similar situations found in many body systems. In fact we have described a gradual modification of the quenched configuration through a procedure that is entirely similar to an annealing protocol. It turns out that a quenched state can gradually be ''melted'' into a state of equilibrium if all the points that have been gained by the agents are periodically eliminated. The points accumulated by the agents keep memory of the previous results and each resetting can be thought of as an exchange of ''energy'' with a thermal bath in order to reach equilibrium.

The concept of temperature has been introduced in the MG in several contexts. In all cases it causes random fluctuations. In this work we have not attempted to introduce the

temperature. It is worthwhile to mention, however, that in fact there is an individual source of fluctuations present, associated with the update of p_i described in Sec. II. We recall that every time a new value of p_i has to be chosen, it is randomly selected from an interval of width δp around it. Hence, this fluctuation acts as a source of fixed temperature, T_0 say, and can be recognized as the final reason that explains why the distribution function $P(p)$ does not simply converge to the sum of several δ's. Hence, our work can be considered to describe the relaxation to equilibrium at that fixed temperature $T = 0$, except when the annealing protocol is applied. This framework could be extended to a situation at $T \neq 0$ by introducing fluctuations, allowing, for instance, individual agents to change their assistance strategies probabilistically.

The BAM has been worked out to show how internal coordination of markets or economic systems can derive from uncoordinated individual actions that are inspired by a common belief. It is a clear example in which an ensemble of agents with limited information about somebody else's decisions can nevertheless give rise to a self-organized pattern within a scheme referred to as ''bounded rationality.'' The relaxation process recalls what in economics is known as ''tatonnement.'' As opposed to this, the ''rational expectations'' solution to the BAM or the GMG would be a distribution $P(p) = \delta(p - \mu)$ because it describes the only common action that is consistent with the general *knowledge* of all agents about everybody else's decisions concerning the acceptable attendance. It is interesting to note that this solution is in general far from equilibrium and hence is neither stationary nor optimal.

The minority game, regarded as a simplified version of the BAM, has also been repeatedly presented as a schematic framework with similar interest for economics. However, the occurrence of quenched asymptotic configurations in a straightforward generalization of the MG puts some limitations on the applicability of the model to stylized economic situations. Quenched configurations are trivial examples of equilibria because the adaptation dynamic has stopped, but they are neither Nash equilibria nor Pareto optimal. In fact, if $M \neq \mu$ many individual situations can be improved without harming the situation of any other agent. It is interesting to note that such configurations are produced by the limited use that agents make of the available information. The GMG can therefore safely be used if quenching is avoided or special situations are considered in which this is not particularly relevant. Otherwise, one should bear in mind that to simulate market environments it is mandatory to assume that the agents must make full use of all the available information.

[1] B. Arthur, Am. Econ. Assoc. Papers Proc. **84**, 406 (1994).

[2] Extensive references on this subject can be found at the web site http://www.unifr.ch/econophysics

[3] D. Challet and Y. C. Zhang, Physica A **246**, 407 (1997); **256**, 514 (1998).

[4] R. Savit, R. Manuca, and R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999); R. Savit, *ibid.* **84**, 1059 (2000).

[5] A. Cavagna, Phys. Rev. E **59**, R3783 (1999); A. Cavagna, J. P. Garrahan, I. Giardina, and D. Sherrington, Phys. Rev. Lett. **83**, 4429 (1999); A. Cavagna, *ibid.* **84**, 1058 (2000).

[6] N. F. Johnson, P. M. Hui, R. Jonson, and T. S. Lo, Phys. Rev. Lett. **82**, 3360 (1999).

[7] E. Burgos and H. Ceva, Physica A **284**, 489 (2000).

[8] E. Burgos, H. Ceva, and R. Perazzo, Physica A **294**, 539 (2001).

[9] D. Heymann, R. P. J. Perazzo, and A. Schuschny, report (unpublished).

[10] A. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, Reading, MA, 1989).

[11] N. F. Johnson, D. J. T. Leonard, P. M. Hui, and T. S. Lo, e-print cond-mat/9905039 v2. See also N. F. Johnson, P. M. Hui, Dafang Zeng, and C. W. Tai, Physica A **269**, 493 (1999).

[12] Within the rules for the BAM, an agent who loses credits is in a position to use the fact that her losses are due to the repeated action of not attending the bar. She could then use this to update her attendance strategy, increasing that probability. This speeds up the convergence to the accepted value of μ in a dramatic way, but adds no conceptual change to the structure of the asymptotic behavior of the system. We did not use this in order to stress that the different behaviors are due only to the different use of the available information.

[13] Note that $\mathcal{C}$ differs from σ_M introduced previously, because the former is the dispersion at time t with respect to the *allowed attendance* $N\mu$, while the latter is the dispersion of the *asymptotic* attendance with respect to its average.