PHYSICAL REVIEW E VOLUME 60, NUMBER 6 DECEMBER 1999

Phase transition and symmetry breaking in the minority game

Damien Challet[1] and Matteo Marsili[2]
[1]*Institut de Physique Théorique, Université de Fribourg, CH-1700 Fribourg, Switzerland*
[2]*Istituto Nazionale per la Fisica della Materia, Unitá di Trieste SISSA, V. Beirut 2-4, Trieste I-34014, Italy*

(Received 30 June 1999)

We show that the minority game, a model of interacting heterogeneous agents, can be described as a spin system and displays a phase transition between a symmetric phase and a symmetry broken phase where the game's outcome is predictable. As a result a "spontaneous magnetization" arises in the spin formalism. [S1063-651X(99)50912-4]

PACS number(s): 02.50.Le, 05.40.−a, 64.60.Ak, 89.90.+n

Market interactions among economic agents give rise to fluctuation phenomena, which are raising much interest in statistical physics [1,2]. The search for a toy system to study agents with marketlike interactions has led to the definition of the minority game (MG) [2,3], a model inspired by Arthur's "El Farol" problem [4], which embodies some basic market mechanisms [2] while keeping the mathematical complexity to a minimum.

In short, the MG is a repeated game where N agents have to decide which of two actions (such as buy or sell) to make. With N odd, this procedure identifies a *minority action* as that chosen by the minority. Agents who took the minority action are rewarded by one payoff unit, whereas the majority of agents looses one unit. Agents do not communicate one with the other and they have access to a "public information," related to past game outcomes, represented by one of P possible patterns.

The strategic point of view of game theory may require, in a case like this, a prohibitive computational task for each of the agents [5]. That is specially true if N and P are very large and agents have no complete information on the detailed mechanism that determines their payoffs, the identity of their opponents, or even their number N. In such complex strategic situations, which are similar to those that agents face in stock markets [2,6], agents may prefer to simplify their decision task by looking for simple behavioral rules that prescribe an action for each of the P possible patterns. This may be particularly advantageous if computational costs exist.

This behavior, called *inductive reasoning* in Ref. [4], is the basis of the MG [2,3]: each agent has a pool of S rules which prescribe an action for each of the P patterns. At each time, she follows her best rule (see below for a more precise definition). These rules, called strategies below, are initially drawn at random among all possible rules, independently for each agent in order to model agents' heterogeneity of beliefs and behaviors.

Numerical simulations [3,7,8] have shown that this system displays a *cooperative* phase for large values of the ratio $\alpha=P/N$: With respect to the simple "random agent" state, where each agent just tosses a coin to choose her action, agents are better off because they get to enstablish a sort of coordination. For small values of α agents receive, on average, poorer payoffs than in the random agent state, a behavior that has been related to crowd effects in markets [2,7,8]. A qualitative understanding of this behavior has been given in terms of geometric considerations [2,9].

In this Rapid Communication we show that the model can be described as a spin system and, as $\alpha=P/N$ varies, it undergoes a dynamical phase transition with symmetry breaking. The symmetry that gets broken is the equivalence between the two actions: in the symmetric phase ($\alpha<\alpha_c$) both actions are taken by the minority with the same frequency (e.g., there are, on average, as many buyers as sellers). For $\alpha>\alpha_c$, in each of the P possible states, the minority does more frequently an action than the other one, i.e., the game's outcome is asymmetric. An asymmetry in the game's outcome is an opportunity that an agent could in principle exploit to gain. This is called an *arbitrage* in economics and it bears a particularly relevant meaning (see discussions in [2,7]). The asymmetry for $\alpha>\alpha_c$ naturally suggests an order parameter and is related to a "phase separation" in the population of agents: while for $\alpha<\alpha_c$ all agents use all of their strategies, for $\alpha>\alpha_c$ a finite fraction ϕ of the agents ends up using only one strategy which, in the spin formalism, is the analog of spontaneous magnetization. The point α_c also marks the transition from persistence (for $\alpha>\alpha_c$) to antipersistence ($\alpha<\alpha_c$) of the game's time series.

Let us start from a sharp definition of the model: We use $+$ and $-$ to denote the two possible actions, so that a generic action is a sign. At each time t, the information available to each agent is the string $\mu_t=(\chi_{t-1},\ldots,\chi_{t-M})$ of the last M actions taken by the minority. This, in our notation, is a string of M *minority signs* $\chi_{t-k}\in\{\pm1\}$. There are $P=2^M$ possible such strings, which we shall label by an index $\mu=1,\ldots,P$ [10]. The index μ_t corresponding to $(\chi_{t-1},\ldots,\chi_{t-M})$ shall be called the present *history*, for short. For each history μ, a strategy a specifies a fixed action a^μ. Each agent $i=1,\ldots,N$ has $S=2$ strategies, denoted by $a_{\pm,i}$, which are randomly drawn from the set of all 2^P possible strategies (the generalization to $S>2$ strategies will be discussed below). We define

$$\omega_i^\mu=\frac{a_{+,i}^\mu+a_{-,i}^\mu}{2},\quad \xi_i^\mu=\frac{a_{+,i}^\mu-a_{-,i}^\mu}{2}$$

so that the strategies of agent i can be written as $a_{s_i,i}^\mu=\omega_i^\mu+s_i\xi_i^\mu$ with $s_i=\pm1$. If $\omega_i^\mu\neq0$, then $\xi_i^\mu=0$ (and vice versa) and the player always takes the decision ω_i^μ whenever the

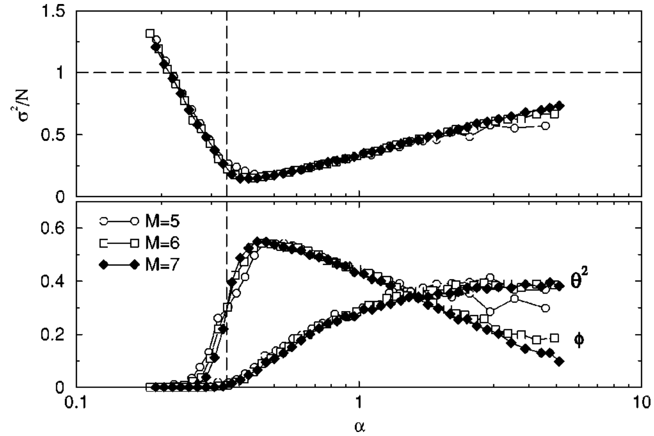

FIG. 1. Top: σ^2/N vs $\alpha = P/N$ for $P = 2^M$ with $M = 5$, 6, and 7. Bottom: θ^2 and ϕ versus α for the same system sizes P. The vertical dashed line is at $\alpha = 0.34 \approx \alpha_c$.

history is μ. The current best strategy of agent i, which she shall adopt at time t, is that which has the highest cumulated payoff. Let us define $\Delta_{i,t} \equiv U_{i,t}^{(+)} - U_{i,t}^{(-)}$ as the difference between the cumulated payoffs $U_{i,t}^{(\pm)}$ of strategies + and − for agent i at time t. Therefore, her choice is given by

$$s_i = \text{sgn}\,\Delta_{i,t}, \tag{1}$$

where ties ($\Delta_{i,t}=0$) are broken by coin tossing. The difference in the population of agents choosing the + and the − sign, at time t, is then

$$A_t = \sum_{i=1}^{N} a_{s_i,i}^{\mu_t} = \Omega^{\mu_t} + \sum_{i=1}^{N} \xi_i^{\mu_t} s_i, \tag{2}$$

where $\Omega^\mu = \sum_i \omega_i^\mu$. The sign chosen by the minority gives the *minority sign* at time t

$$\chi_t = -\text{sgn}\,A_t \tag{3}$$

and this determines the new history μ_{t+1}, which corresponds to the string $(\chi_t, \ldots, \chi_{t-M+1})$ [10]. Finally, each agent i rewards those of her strategies which have predicted the right sign ($a_{s,i}^{\mu_t} = \chi_t$) updating the cumulated payoffs $U_{i,t+1}^{(\pm)} = U_{i,t}^{(\pm)} + a_{\pm,i}^{\mu_t}\chi_t$. This implies that the cumulated payoff difference $\Delta_{i,t}$ is updated according to

$$\Delta_{i,t+1} = \Delta_{i,t} + 2\chi_t \xi_i^{\mu_t}. \tag{4}$$

Equations (1)–(4) update the state $\{\mu_t, \Delta_{i,t}\}$ of the system from t to $t+1$. With an initial condition (e.g., $\mu_0 = 1$, $\Delta_{i,0} = 0$, $\forall i$) the dynamics of the MG is completely specified. The "quenched" variables $\{\Omega^\mu, \xi_i^\mu\}$ play here the same role as disorder in statistical mechanics [11].

An important quantity in the MG is the variance $\sigma^2 = \langle A^2 \rangle$ of the difference A in the sizes of the two populations, where $\langle \cdot \rangle$ is a time average in the stationary state of the process specified by Eqs. (1)–(4). The number of winners, at each time step, is $(N - |A|)/2 \approx (N-\sigma)/2$ so that smaller fluctuations σ^2 correspond to larger global gain. A population of random agents would yield $\sigma^2 = N$. Numerical simulations [3,7,8] (see Fig. 1) show that, for $\alpha = P/N$ large enough, agents with inductive reasoning manage to behave globally better (i.e., $\sigma^2 < N$) than random agents, whereas $\sigma^2 > N$ for small α (see Fig. 1). However, no singularity (and no order parameter) has been yet identified in order to locate a phase transition.

As shown in Ref. [12], to a good approximation one can neglect the coupling of the dynamics of $\Delta_{i,t}$ and μ_t and replace the dynamics of the latter by random sampling of the history space, i.e., Prob$(\mu_t = \mu) = 1/P$, $\forall \mu$. This simplifies considerably our discussion since then

$$\sigma^2 \simeq \frac{1}{P}\sum_{\mu=1}^{P} (\Omega^\mu)^2 + 2\sum_{i=1}^{N} h_i \langle s_i \rangle + \sum_{i,j=1}^{N} J_{i,j} \langle s_i s_j \rangle, \tag{5}$$

where $\langle \cdot \rangle$ stands for a time average and

$$h_i = \frac{1}{P}\sum_{\mu=1}^{P} \Omega^\mu \xi_i^\mu, \quad J_{i,j} = \frac{1}{P}\sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu. \tag{6}$$

The field h_i measures the difference of correlation of the two strategies with Ω^μ, whereas the coupling $J_{i,j}$ accounts for the interaction between agents as well as for agents self-interaction ($J_{i,i}$). The structure of the couplings (6) is reminiscent of neural networks models [11], where ξ_i^μ play the role of memory patterns. This similarity confirms the conclusion of Refs. [2,7,9] that the relevant parameter is the ratio $\alpha = P/N$ between the number of patterns and the number of spins.

The key element that is at the origin of the behavior of the model is the fact that for each history μ, there are agents which always take the same decision. This gives rise to the time independent contribution Ω^μ in A, which produces a bias in the value of χ_t whenever $\mu_t = \mu$. A measure of this bias, is given by the parameter

$$\theta = \sqrt{\frac{1}{P}\sum_{\mu=1}^{P} \langle \chi | \mu \rangle^2}, \tag{7}$$

where $\langle \chi | \mu \rangle$ is the conditional average of χ_t given that $\mu_t = \mu$. Loosely speaking, θ measures the presence of information or arbitrages in the signal χ_t. If $\theta > 0$ an agent with strategies of "length" $M = \log_2 P$ can detect and exploit this information if one of her strategies is more correlated with $\langle \chi | \mu \rangle$ than the other. More precisely, we observe that if $v_i \equiv \langle \Delta_{i,t+1} - \Delta_{i,t} \rangle \neq 0$, then $\Delta_{i,t} \simeq v_i t$ grows linearly with time, and the agent's spin will always take the value $s_i = \text{sign}\,v_i$. We shall call this a *frozen* agent, since her spin variable is frozen. We find

$$v_i = \langle \chi_t \xi_i^{\mu_t} \rangle \simeq \frac{1}{P}\sum_{\mu=1}^{P} \langle \chi | \mu \rangle \xi_i^\mu \propto -h_i - \sum_{j=1}^{N} J_{i,j} \langle s_j \rangle, \tag{8}$$

where the last equation relies on an expansion of $\langle \chi | \mu \rangle$ to linear order in A [13].

It is instructive to consider first the case where other agents choose by coin tossing (i.e., $\langle s_j \rangle = 0$ for $j \neq i$) so that $v_i \propto -h_i - J_{i,i}\langle s_i \rangle$. If $v_i \neq 0$, then $s_i = \text{sgn}\,v_i = -\text{sgn}(h_i + J_{i,i}\langle s_i \rangle)$. But this last equation has a solution only if $|h_i| > J_{i,i}$, whereas otherwise $|\langle s_i \rangle| < 1$ and $v_i = 0$. Note that $J_{i,i} \simeq 1/2$ and that h_i can be approximated by a Gaussian variable with zero average and variance $(4\alpha)^{-1}$. This means that $|h_i| \ll J_{i,i}$ for $\alpha \gg 1$, which implies that most agents have $\langle s_i \rangle \approx 0$ in this limit and we can indeed neglect agent-agent interaction. This allows one to compute the probability for an agent to be frozen,

$$\phi = P\{|h_i| > J_{i,i}\} \propto e^{-\alpha/2}, \tag{9}$$

for $\alpha \gg 1$. Numerical simulations show that $\phi \propto e^{-(0.37\pm0.02)\alpha}$ indeed decays exponentially. As $\alpha \to \infty$, the random agents limit is attained because $\langle s_i \rangle \to 0$ for all i and $\langle s_i s_j \rangle = \langle s_i \rangle \langle s_j \rangle$ for $i \neq j$. By Eq. (5) we find $\sigma^2 = \Sigma_\mu (\Omega^\mu)^2/P + \Sigma_i J_{i,i} \simeq N$.

The same argument applies in general, with the difference that the "bare" field h_i must be replaced by the "effective" field $\tilde{h}_i = h_i + \Sigma_{j\neq i} J_{i,j} \langle s_j \rangle$. In order for agent i to get frozen, her effective field $\tilde{h}_i$ must ovecome the self-interaction $J_{i,i}$, i.e., $|\tilde{h}_i| > J_{i,i} \simeq 1/2$. If this condition is met, $s_i = -\operatorname{sgn} \tilde{h}_i$. It can also be shown that a frozen agent will, on average, receive a larger payoff than an unfrozen agent [14]. Loosely speaking, one can say that a frozen agent has a *good* and a *bad* strategy and the good one remains better than the bad one even when she actually uses it. On the contrary, unfrozen agents have two strategies, each of which seems better than the other when it is not adopted. In this sense, symmetry breaking in $\langle \chi | \mu \rangle$ induced a sort of breakdown in the *a priori* equivalence of agents' strategies.

A quantitative analysis of the fully interacting system shall be presented elsewhere [14]. For the time being we shall discuss the behavior of the system on the basis of extensive numerical simulations. Figure 1 reports the behavior of θ, ϕ and σ^2 as functions of α for several values of P. As α decreases, i.e., as more and more agents join the game, the arbitrages opportunities, as measured by θ decrease. In loose words, agents' exploitation of the signal Ω^μ weakens its strength by screening it with their adaptive behavior. If the number N of agents is small compared to the signal "complexity" $P = 2^M$, agents exploit only partially the signal Ω^μ, whereas if $N \gg P$ then Ω^μ is completely screened by agents' behavior and $\theta = 0$. As Figure 1 shows the parameter θ displays the characteristic behavior of an order parameter with a singularity at $\alpha_c \approx 0.34$. Accordingly, also the fraction ϕ of frozen agents drops to zero as $\alpha \to \alpha_c^+$. The comparison between different system sizes in Fig. 1 strongly suggests that ϕ drops discontinuously to zero at α_c (and it also gives the value of α_c). The vanishing of ϕ is clearly a consequence of the fact that θ also vanishes at α_c. Indeed if $\langle \chi | \mu \rangle = 0$ for all μ, by Eq. (8), also $v_i = 0$ for all i, so that $\Delta_{i,t}$ remains bounded and $|\langle s_i \rangle| < 1$.

The transition can also be understood in terms of the variables $\Delta_{i,t}$ as an "unbinding" transition as $\alpha \to \alpha_c^-$: For $\alpha < \alpha_c$ a "bound state" exists with finite $\Delta_{i,t}$, which corresponds to the fact that the equations $v_i = 0$, $i = 1, \ldots, N$ admit a solution with $|\langle s_i \rangle| < 1$, $\forall i$ [14] (only P of the equations $v_i = 0$ are linearly independent). For $\alpha > \alpha_c$ this is no longer true and the population separates: a fraction ϕ of variables $\Delta_{i,t}$ acquire a constant "velocity" $v_i \neq 0$ (with $|\langle s_i \rangle| = 1$), whereas for the remaining agents $v_i = 0$, $\Delta_{i,t}$ remains bounded and $|\langle s_i \rangle| < 1$.

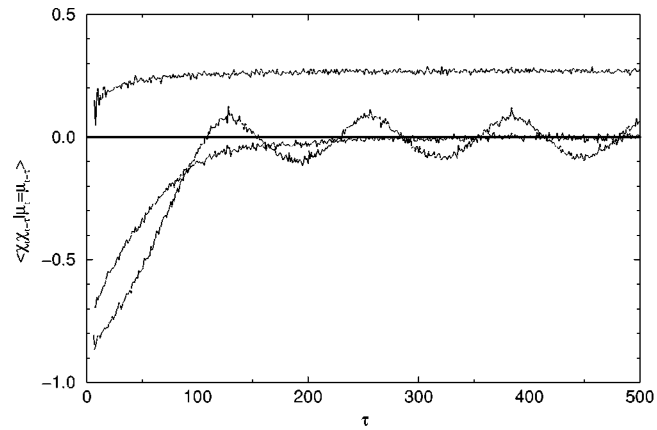

FIG. 2. Temporal correlation of χ_t on the same history, $\langle \chi_t \chi_{t+\tau} | \mu_t = \mu_{t+\tau} \rangle$, averaged over all histories vs τ (10^6 iterations, $M=6$, $\alpha = 0.5$, 0.22, and 0.1)

It is suggestive to observe that $v_i \propto -(\partial \sigma^2/\partial s_i)$ so that the dynamics of the minority game is actually similar to a spin dynamics with Hamiltonian σ^2. Indeed either the spin is frozen in the direction that minimizes $-s_i v_i(s_i)$, or its average $\langle s_i \rangle$ is such that $v_i = 0$. This then explains why cooperation occurs in the MG. A closer analysis, to be reported elsewhere [14], reveals that indeed the stationary state of the MG is described by the ground state properties of an Hamiltonian very similar to σ^2. Finite size scaling suggests that σ^2 has a minimum at α_c with a discontinuity in its derivative (see Fig. 1). These conclusions are indeed confirmed by exact results [14]. It is worth stressing, however, that the qualitative aspects of the transition are already captured at the simple level of approximation of Eq. (8).

Let us go back to Fig. 1. Above α_c agents do not fully exploit the information Ω^μ and, as a result, $\langle \chi | \mu \rangle \neq 0$. Figure 2 shows that χ_t shows persistence in time, in the sense that when $\mu_t = \mu_{t+\tau}$ the minority signs χ_t and $\chi_{t+\tau}$ tend to be the same. This persistence disappears, $\langle \chi_t \chi_{t+\tau} | \mu_t = \mu_{t+\tau} \rangle \to 0$ as α decreases and it turns into anti-persistence for smaller α. The oscillatory behavior in Fig. 2 has indeed period $2P$, which means that typically when the population comes again on the same history μ it tends to do the opposite of what it did the time before. Even if finite size effects do not allow a definite conclusion, it is quite likely that this change in time correlations also occurs at α_c [14]. Time correlations, even though of opposite nature, are present both above and below α_c. These are like arbitrages in a market which could be exploited by agents. In this sense the market is *efficient*, i.e., arbitrage free, only for $\alpha = \alpha_c$.

The same qualitative behavior is expected when agents have $S > 2$ strategies. Again for a given history μ it may happen that all the S agent's i strategies prescribe the same action: agent i will do that action no matter what strategy she has chosen. As S increases, this will occur for a smaller and smaller number of histories (more precisely with a probability 2^{1-S}). This shall correspond to a weaker signal Ω^μ, which is in complete agreement with the observation [7,9] of shallower features for larger S. Note that for each agent it would be rewarding to increase the number of strategies because they would have more chances to outguess χ_t. At the same time, if all agents increase S the game becomes less rewarding for all of them, at least for $\alpha > \alpha_c$. This situation is typical of games, such as the *tragedy of commons*, where many agents interact through a global resource [15].

The condition $v_i=0$ for the bound state in the symmetric phase involves P equations with $(S-1)N$ variables. This suggests that in general the scaling parameter is $\alpha=P/[(S-1)N]$. The curve σ^2/N as a function of $\alpha=P/[(S-1)N]$ collapse remarkably well one on the other for $\alpha\leqslant\alpha_c$ (especially for $S>2$) but not for $\alpha>\alpha_c$ [e.g., in the large α behavior $\phi\propto e^{-C(S)\alpha}$ we found $C(2)\approx 0.37$, $C(3)\approx 1.50$ and $C(4)\approx 2.90$].

Our approach also implies that no coordination is possible if agents have $S=2$ opposite strategies ($a^{\mu}_{+,i}=-a^{\mu}_{-,i}$) because then $\Omega^{\mu}=0$. Numerical simulations show that indeed $\sigma^2\geqslant N$ for all $\alpha>0$ in this case.

The same qualitative behavior also occurs in a wide range of related models. First, total freezing occurs in majority models. Note indeed that changing the sign of Eq. (3) would also change the sign in Eq. (8). In particular the self-interaction $J_{i,i}$ changes sign so that it becomes favorable for each agent to stick to only one strategy anyway. The model is therefore trivial. More interesting models are obtained keeping the ''frustration'' effects of the MG but changing the definition of payoffs in Eq. (4). It can be shown [14] that the phase transition and the large α behavior are quite robust features of minority games (see, e.g., [13]).

In summary, we find that a phase transition occurs in the minority game. The cooperative phase ($\alpha>\alpha_c$) is characterized by the presence of a fraction ϕ of frozen agents (who use only one strategy), unexploited arbitrages ($\langle\chi|\mu\rangle\neq 0$), and persistence in the global signal χ_t. In the symmetric phase ($\alpha<\alpha_c$) inductive dynamics is inefficient: agents adopt strategies when they are no longer good. There is no arbitrage (for strategies of length M) to exploit and the signal shows anti-persistence.

We acknowledge Y.-C. Zhang for enlightening discussions, useful suggestions, and for introducing us to the minority game. This work was partially supported by Swiss National Science Foundation Grant No. 20-46918.98.

[1] P. W. Anderson, K. Arrow, and D. Pines, *The Economy as an Evolving Complex System* (Addison-Wesley, New York, 1988); *Econophysics: an Emerging Science*, edited by J. Kertesz and I. Kondor (Kluwer, Dordrecht, 1998).

[2] Y.-C. Zhang, Europhys. News **29**, 51 (1998).

[3] D. Challet and Y.-C. Zhang, Physica A **246**, 407 (1997).

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

[5] A strategic player has to consider, for any given strategy she can actually play, how each of her opponents would react knowing that she is playing that strategy, and how that affects her payoff. On this basis, which assumes that also her opponents behave the same, she has to find out what the best strategy is. See, e.g., D. Fudenberg and J. Tirole, *Game Theory* (MIT Press, Cambridge, MA, 1991).

[6] Y.-C. Zhang, e-print cond-mat/9901243.

[7] R. Savit, R. Manuca, and R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999); e-print adap-org/9811005.

[8] N. F. Johnson *et al.*, Physica A **256**, 230 (1998).

[9] D. Challet and Y.-C. Zhang, Physica A **256**, 514 (1998); N. F. Johnson *et al.*, e-print cond-mat/9811227.

[10] An example of labeling is binary coding. μ_{t+1} is then given by: $\mu_{t+1}=[2\mu_t+1+\chi_t/2]\mathrm{mod}\,P$.

[11] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).

[12] A. Cavagna, Phys. Rev. E **59**, R3787 (1999) has shown that replacing the dynamics of μ_t by random sampling does not change, to a very good approximation, the behavior of the model. The assumption that each history μ is sampled with the same frequency is strictly true only in the symmetric phase $\alpha<\alpha_c$ but it holds qualitatively also for $\alpha>\alpha_c$.

[13] This approximation is justified by the fact that the ''linear'' minority model obtained replacing Eq. (4) by $\Delta_{i,t+1}=\Delta_{i,t}-A_t\xi_i^{\mu_t}$ gives results very similar to those presented here. An analytic derivation is also possible from the identity $\langle\chi|\mu\rangle=2\,\mathrm{Prob}(A^{\mu}<0|\mu)-1$ [14].

[14] D. Challet, M. Marsili, and R. Zecchina (unpublished).

[15] G. Hardin, Science **162**, 1243 (1968).