ELSEVIER
Physica A 314 (2002) 83–91
PHYSICA A
www.elsevier.com/locate/physa

Stochastic decision-making in the minority game

D. Sherrington[a,*], A.C.C. Coolen[b], J.A.F. Heimel[b]

[a] *Department of Physics, University of Oxford, Theoretical Physics, 1 Keble Road, Oxford, OX1 3NP, UK*
[b] *Department of Mathematics, King's College London, The Strand, London WC2R 2LS, UK*

Abstract

A discussion is presented of the effects of stochasticity in the decision-making of agents in the minority game. Both simulational and analytic results are reported and discussed for both additive and multiplicative noise. As a function of the ratio d of information dimension to number of agents a phase transition separates a low d non-ergodic phase from a high d ergodic phase. For additive noise the critical d_c is temperature-independent but for multiplicative noise $d_c(T)$ decreases with T. Additive noise does not affect the asymptotic behaviour for $d > d_c$ but is relevant below d_c. Multiplicative noise has consequence for all d.

PACS: 02.50.Le; 87.23.Ge; 05.70.Ln; 64.60.Ht

Keywords: Econophysics; Statistical physics; Minority game

The Minority Game (MG) [1] is a simple minimalist model inspired by considerations of a stock market of speculators attempting to profit by buying low and selling high, basing their decisions on the application of individual strategies to common information and learning from experience. Aside from its potential interest from the perspective of economics, the MG has fundamental interest from the viewpoint of the statistical physics of complex systems and it is this latter which is the emphasis of the present paper, with particular reference to the role of stochastic decision-making. After a general introduction, we discuss briefly simulational results for the stochastic case and then present results of an exact analytic solution.

In the MG the speculators are replaced by 'agents' who, at each time step make either of two choices, with the individual objectives of making the opposite choice to the majority of agents. To determine their choices each agent has a small set of quenched

* Corresponding author. Tel.: +44-1865-27396; fax: +44-1865-273947.
E-mail address: d.sherrington1@physics.ox.ac.uk (D. Sherrington).

0378-4371/02/$ - see front matter
PII: S0378-4371(02)01158-5

randomly and independently chosen strategies whose application to the current common information yields a choice. The strategy actually employed by an agent at any step is decided on the basis of the current value of 'points' awarded cumulatively so as to 'reward' strategies whose application would have yielded the actual minority choice; this provides the learning. Individual agents have no direct knowledge of one another and there are no other rewards. The system thus has the key ingredients of frustration and quenched disorder which underly the complex cooperative behaviour of several more conventional physical systems. It is also mean-field like, thus offering the potential of exact solution, which we indicate below.

In the original model [1] the information was the minority choice for the last m time-steps, the strategies were Boolean functions operating on this information and those used were selected deterministically, each agent employing his/her strategy with the currently highest point-score. In fact, however, the qualitative macroscopic behaviour is essentially unaffected by replacing the historical information by a random, but common, choice at each time step [1] [2] and by replacing the space of 2^m information options by a D-dimensional vector space, with the strategies and the information both D-dimensional vectors and the choice/bid given by their scalar product [3]. We employ this modification below.

The most natural instructive macroscopic observable is the volatility σ, the variance of the total bid. It exhibits several interesting features as shown initially in simulations; (i) it deviates from the value which would result from agents acting randomly [4], thereby demonstrating correlation, (ii) it shows scaling, $\sigma/\sqrt{N}$ versus $d = D/N$ where N is the number of agents [4]; (iii) for a start from an unbiased point-score, as a function of d it varies from a value greater than random for low d, reduces to a minimum below the random value at some $d_m \sim \mathcal{O}(1)$, and approaches the random value asymptotically as $d \to \infty$ [5]; (iv) for a start from a randomly strongly biased point-score the behaviour is identical to that for the unbiased case for d greater than a critical value d_c, with d_c close to d_m, suggesting ergodicity, but for $d < d_c$ there is non-ergodicity and σ continues to decrease as d is lowered through d_c [6].

The higher than random volatility for the case of unbiased start for low d is due to crowding effects [7] and can be reduced by adding stochasticity to the agents' decision-making. It might seem natural to expect that σ could be reduced to its random value, but in fact it can be reduced further [3] and the behaviour is actually more subtle.

Stochasticity can be added to the agents' decision-making by replacing the original deterministic rule of employing their strategies with the highest current point-scores by probabilistic ones. Since the above qualitative features are maintained for any finite number $s > 1$ of strategies per agent, for simplicity, we concentrate on $s=2$ where only the difference $p_i(t)=P_i^1(t)-P_i^2(t)$ between the point-scores of an agent's two strategies is relevant; here i labels the agent. In keeping with the minimalist philosophy it is natural to consider first the case of a single control parameter for the stochasticity. Thus we might take the relative probability that agent i uses his/her first, rather than second, strategy to be given by $\exp(f(p_i(t))/T)$. For all forms of $f(p)$ with $\mathrm{sgn}[f(p)]=\mathrm{sgn}[p]$

[1] This reflects the fact that the 'information' principally provides an effective interaction 'field' between the agents.

that have been studied, and presumably for all such $f(p)$, $\sigma(T) \leqslant \sigma(0)$ and $\mathrm{Min}_T\, \sigma(T) \leqslant \sigma_{\rm random}$. However, two classes of function $f(p)$ can be identified and have distinct features. One is unbounded as $|p|$ grows and has its simplest manifestation in the form [3]

$$f(p) = p \tag{1}$$

while the second is bounded, with its simplest form [6]

$$f(p) = \mathrm{sgn}[p] ; \tag{2}$$

for reasons which will become clearer later these will be referred to, respectively, as additive and multiplicative noise. The reason for the need for the distinction lies in the phenomenon of frozen agents [8] whose coarse-time-averaged point-scores $p_i(t)$ maintain their sign and grow quasi-linearly in magnitude. For large enough times such agents effectively behave as deterministic for any T for additive noise (1) [9], but become random as $T \to \infty$ for multiplicative noise (2) [6]. Again there is a distinction in behaviour for d greater or less than d_c. As d_c is decreased from ∞ the fraction of frozen agents increases monotonically. For the *tabula rasa* start it then drops rapidly to zero around d_c but continues to increase for the biased start. For additive noise (1), for $d > d_c$ the volatility is T-independent provided one waits long enough [9], but for $d < d_c$ a sufficient $T > T_c(d)$ leads to a σ which converges to a value intermediate between the results of deterministic ($T = 0$) rules for *tabula rasa* and biased starts [11,12], corresponding to the minimization of an effective Hamiltonian [10]. For multiplicative noise, at any $d < d_c$ there is a critical $T(d)$ which corresponds to a minimum of $\sigma(T)$, but for $d > d_c$ any $T > 0$ increases the volatility (but it is still bounded above by $\sigma_{\rm random}$).

The above is based on computer simulations and iterations of derived microscopic equations. Let us now turn to analytic considerations, concentrating here on an exact solution of the macroscopic dynamics via the medium of generating functional theory [13,14] and on batch dynamics in which the original procedure of updating the points every microscopic time step in response to instantaneous random information (online dynamics) is replaced by an update (at a re-scaled macroscopic time) corresponding to an average over possible information. [2] Taking the information $\boldsymbol{I}(t)$ as a stochastic unit vector in D-dimensional space and the strategies $\boldsymbol{R}_i^\alpha$; $\alpha = 1, 2$ as quenched random unit-vectors in the same space, [3] the dynamical update becomes [16]

$$p_i(t+1) = p_i(t) - h_i - \sum_j J_{ij} s_j(t) , \tag{3}$$

where

$$h_i = N^{-1} \sum_i \boldsymbol{\xi}_i \cdot \boldsymbol{\omega}_j , \tag{4}$$

[2] The batch dynamics is equivalent to carrying out the usual procedure of response to random information but updating the points only after a number of steps $\geqslant \mathcal{O}(D)$, and was shown to produce identical stationary order parameters as the original on-line model for $d > d_c$ [15].

[3] Variously the D-dimensional space has been taken as hyperspherical or hypercubical. Bold symbols denote vectors in these spaces.

$$J_{ij} = N^{-1}\boldsymbol{\xi}_i \cdot \boldsymbol{\xi}_j \; , \tag{5}$$

$$\boldsymbol{\xi}_i = \tfrac{1}{2}(\boldsymbol{R}_i^1 - \boldsymbol{R}_i^2) \; , \tag{6a}$$

$$\boldsymbol{\omega}_i = \tfrac{1}{2}(\boldsymbol{R}_i^1 + \boldsymbol{R}_i^2) \tag{6b}$$

and $s_j(t)$ is a stochastically determined 'spin' related to $p_j(t)$ via

$$s_j(t) = \sigma[p_j(t), z_j(t)|T] \tag{7}$$

with

$$\sigma[p,z|T] = \mathrm{sgn}[p+Tz]; \quad \text{additive noise} \; , \tag{8a}$$

$$\sigma[p,z|T] = \mathrm{sgn}[p]\,\mathrm{sgn}[1+Tz]; \quad \text{multiplicative noise} \tag{8b}$$

and the $z_j(t)$ are zero-average random numbers chosen independently at each time step from a normalized distribution $P(z)$ with unit variance. [4] We now take (3) as the defining equation for the dynamics, with the addition to h_i of perturbation fields $\theta_i(t)$, used to generate response functions but taken to zero at the end.

The generating functional is then

$$Z = \int \prod_t [\mathrm{d}\vec{p}(t) W(\vec{p}(t+1)|\vec{p}(t))] P_0(\vec{p}(0)) \exp\left\{ i \sum_t \vec{\Psi}(t) \cdot \vec{p}(t) \right\} \; , \tag{9}$$

where $W(\vec{p}|\vec{p}\,')$ is the transition probability given by (3) and $P_0(\vec{p}(0))$ is the distribution of $\vec{p}(t)$ at $t=0$. [5] $\vec{\Psi}(t)$ is an auxiliary generating field, taken to zero at the end. Averaging over the $\boldsymbol{R}$ yields typical behaviour, while the range-free nature of the problem permits the elimination of microscopic variables in favour of macroscopic correlation and response functions and the further evaluation of dominant behaviour via steepest descent analysis to produce an effective non-Markovian self-consistent single particle dynamics, from which the macroscopic parameters of interest may be obtained [16].

The effective single particle dynamics is [6]

$$p(t+1) = p(t) - d \sum_{t \leqslant t'} (\underline{1} + \underline{G})^{-1}_{tt'} \sigma[p(t'), z(t')|T] + \theta(t) + \sqrt{d}\,\eta(t) \; , \tag{10}$$

where the $z(t)$ are the original single-agent decision noises chosen independently randomly at each time step with probability distribution $P(z)$, the $\theta(t)$ are perturbation fields (again introduced to generate response functions), the $\eta(t)$ provide a second self-consistent source of effective decision noise arising from the inter-agent correlations, and $\underline{G}$ is a self-consistent response function (Fig. 1). The $\eta(t)$ are

[4] Eqs. (1) and (2) result from a simple special case of this distribution.

[5] Overarrows are used to denote vectors in agent-space and are N-dimensional.

[6] Underlined symbols are matrices in time space. Tildes will be employed to denote vectors in this space. Subscripts or bracketed symbols are used interchangeably to denote components.

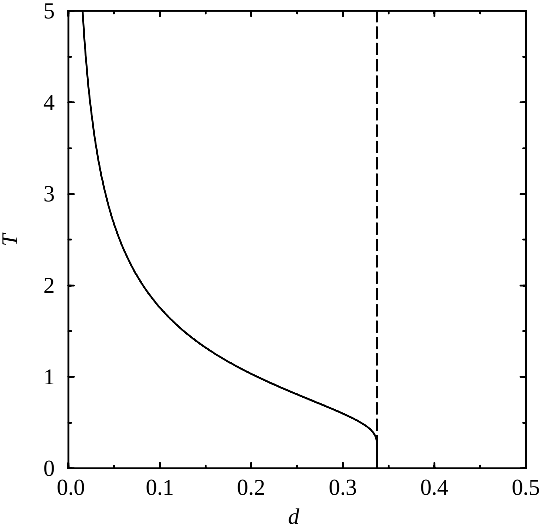

Fig. 1. Phase diagram in the (d, T) plane. The lines separate ergodic phases on the right from non-ergodic phases on the left. (i) solid line: multiplicative noise, (ii) dashed line: additive noise. From Ref. [16].

Gaussian-distributed with zero mean and temporal correlations related to the correlation and response functions by

$$\langle \eta(t)\eta(t')\rangle = \Sigma_{tt'} \; , \tag{11}$$

$$\underline{\Sigma} = (\underline{1} + \underline{G})^{-1}\underline{D}(\underline{1} + \underline{G}^T)^{-1} \; , \tag{12}$$

$$D_{tt'} = 1 + C_{tt'} \; , \tag{13}$$

where

$$C_{tt'} = \langle s(t)s(t')\rangle_* |_{\{\theta(t'')=0\}} \; , \tag{14}$$

$$G_{tt'} = \frac{\partial}{\partial\theta(t')} \langle s(t)\rangle_* |_{\{\theta(t'')=0\}} \; , \tag{15}$$

$$s(t) = \sigma[p(t), z(t)|T] \; , \tag{16}$$

$$\langle f(\tilde{p}, \tilde{z})\rangle_* = \int \prod_t \left[\frac{\mathrm{d}p(t)}{\sqrt{2\pi}}\right] \langle M(\tilde{p}, \tilde{z}) f(\tilde{p}, \tilde{z})\rangle_z \; , \tag{17}$$

$$M(\tilde{p}, \tilde{z}) = P_0(p(0)) \int \prod_t \left[\frac{\mathrm{d}q(t)}{\sqrt{2\pi}}\right] \exp\left\{ -\frac{\mathrm{d}}{2}\, \tilde{q} \cdot (\underline{1} + \underline{G})^{-1} \cdot \underline{D} \cdot (\underline{1} + \underline{G}^T)^{-1} \cdot \tilde{q} \right.$$
$$+ i \sum_t q(t) \left[p(t+1) - p(t) - \theta(t) \right.$$
$$\left.\left. + d \sum_{t'} (\underline{1} + \underline{G})^{-1}_{tt'} \sigma[p(t'), z(t')|T] \right] \right\} \; , \tag{18}$$

and $\langle\,\rangle_z$ denotes an average over z, with measure $P(z)$. Eqs. (10)–(18) then define the closed set of equations describing the macroscopic behaviour. They can be related to more conventional macroscopic measures by identification with the macroscopic autocorrelations and autoresponse functions

$$C_{tt'} = \lim_{N\to\infty} N^{-1} \sum_i \overline{\langle s_i(t) s_i(t')\rangle}|_{\{\theta(t)=0\}} \,, \tag{19}$$

$$G_{tt'} = \lim_{N\to\infty} N^{-1} \sum_i \frac{\partial}{\partial\theta_i(t')} \overline{\langle s_i(t)\rangle}|_{\{\theta(t)=0\}} \,, \tag{20}$$

where $\overline{\langle\,\rangle}$ denotes an average of the many agent dynamics over the choice of strategies.

Eqs. (10)–(18) are exact in the thermodynamic limit $N \to \infty$. They are still complicated in their self-consistency which is highly non-linear and non-local in time. However, already from the simulations it is known that an ergodic solution should be anticipated for $d > d_c(T)$ and so we concentrate first on this region in the asymptotic long time limit in which one reaches a time-translationally invariant stationary state with $G_{tt'} = G(t-t')$ and $C_{tt'} = C(t-t')$. For $d > d_c(t)$ this state has no anomalous response; i.e., $\lim_{\tau\to\infty} \sum_{t\leqslant\tau} G(t) = k$ exists. The lower limit of such behaviour defines $d_c(T)$. The corresponding persistent quantity for the autocorrelation is $\lim_{\tau\to\infty} \sum_{t\leqslant\tau} C(t) = c$ (Fig. 2).

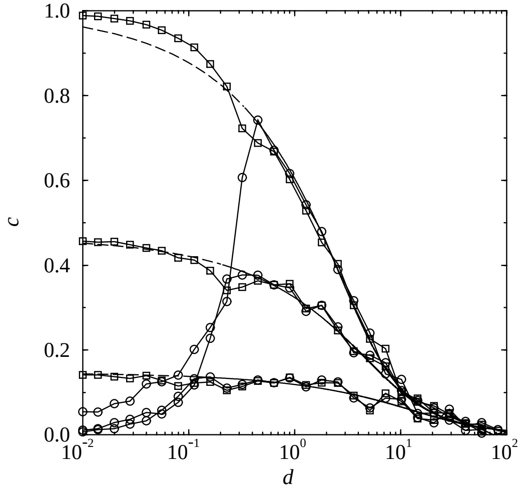

Fig. 2. Persistent correlation c as a function of d for three choices of multiplicative noise (bottom to top: $T = 0, 1, 2$). Thick solid curves: analytic predictions for $d \geqslant d_c(T)$. Thick dashed curves: continuations of (21) beyond its obvious limit of validity. Connected markers: simulation runs: circles $p_i(0) = 0$, squares $|p_i(0)| = 10$. From Ref. [16].

The quantities c and k are then given by

$$c = \varDelta^2 - \left[\varDelta^2 - \frac{(1+c)}{d}\right] \operatorname{erf}\left[\sqrt{\frac{d\varDelta^2}{2(1+c)}}\right] - 2\varDelta\sqrt{\frac{1+c}{2\pi d}} \exp\left\{-\frac{d\varDelta^2}{2(1+c)}\right\} , \tag{21}$$

$$k = \left\{\frac{d}{\operatorname{erf}[\sqrt{\frac{d\varDelta^2}{2(1+c)}}]} - 1\right\}^{-1} , \tag{22}$$

where for additive noise

$$\varDelta = 1 \tag{23a}$$

and for multiplicative noise (Fig. 3)

$$\varDelta = \int \mathrm{d}z\, P(z)\, \operatorname{sgn}[1 + Tz] . \tag{23b}$$

Thus indeed, within the region of validity of the assumption of non-anomaly additive noise has no consequence and the behaviour is exactly as the deterministic case, with $d_c \approx 0.33740$. Both c and k go to zero as $d \to \infty$ (for any T and either type of noise) and grow monotonically as d is reduced, with $k \to \infty$ signalling $d \to d_c$. In the case of multiplicative noise d_c is temperature-dependent, reducing from the deterministic value towards zero as $T \to \infty$. Fig. 1 shows the resultant phase diagram. For $d < d_c(T)$ the system behaves non-ergodically, and the assumptions which led to (21), (22) are no longer obviously valid, but the solution to (21) continues to provide a reasonable fit to simulations starting from a strongly biased state. Fig. 2 illustrates the predictions and comparison with simulation.

The density of frozen agents $\varPhi$ is given by

$$\varPhi = (1 - k(1-d))/(1+k) , \tag{24}$$

which again fits simulations well for $d > d_c$ and also approximately for $d < d_c$ for the case of strongly biased start; see Fig. 3. The volatility (and its corresponding non-temporally local extensions) involves short-term as well as long-term fluctuations, thereby requiring more than just c and k, but can be evaluated with consistent ansätze [16], again comparing well with simulations for $d > d_c$. In this case the continuation below d_c yields results intermediate between those of *tabula rasa* and strongly biased starts; see Fig. 4.

Complete solutions of (10)–(18) for $d < d_c(T)$ remain to be performed, but we note (i) that in the deterministic case both high and low volatility stationary states of the corresponding equations have been demonstrated [14] and shown to be in good accord with the two extremes of initialization, (ii) that in an analysis of the first few time steps from a *tabula rasa* start both additive and multiplicative noise serve to reduce oscillations in the $p_i(t)$ from one time step to the next, thereby destabilizing the high volatility state [16], (iii) continuity arguments suggest that in the analytic

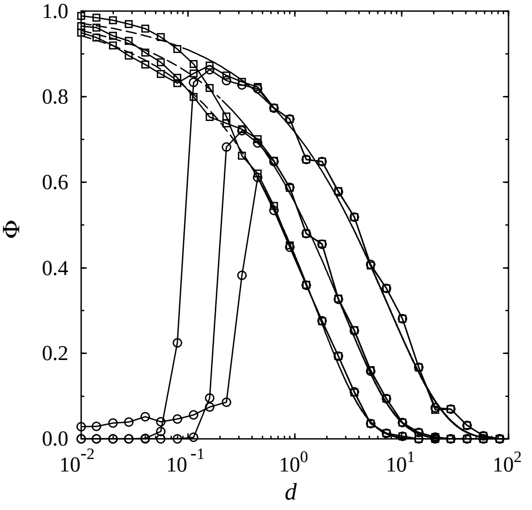

Fig. 3. Asymptotic fraction of frozen agents as a function of d for three choices of multiplicative noise (bottom to top: $T = 0, 1, 2$). Thick solid curves: analytic predictions for $d \geqslant d_c(T)$. Thick dashed curves: continuations beyond validity. Connected markers: simulations: circles $p_i(0) = 0$, squares $|p_i(0)| = 10$. From Ref. [16].

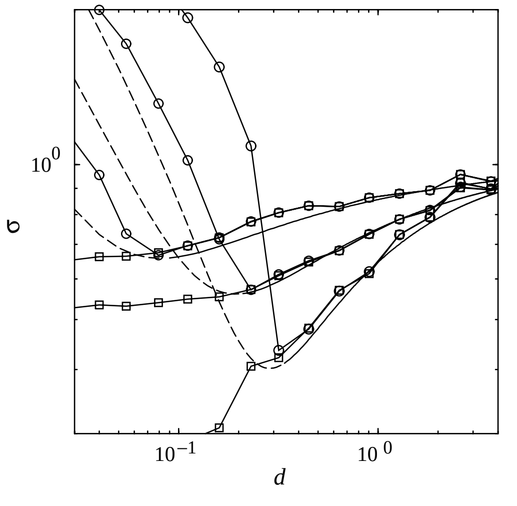

Fig. 4. Asymptotic volatility as a function of d for three choices of multiplicative noise (bottom to top: $T = 0, 1, 2$). Thick solid lines: analytic predictions for $d \geqslant d_c(T)$. Thick dashed lines: continuations beyond obvious validity. Connected markers: simulations: circles $p_i(0) = 0$, squares $|p_i(0)| = 10$. From Ref. [16].

continuation of σ below $d_c(T)$ the minimum $d_m(T) \leqslant d_c(T)$ for multiplicative noise and simulations [16] bear this out.

The analysis of (10)–(17) is readily extended to a non-homogenous ensemble of agents with a distribution of stochastic control parameters; $\langle f(\tilde{p}, \tilde{z}) \rangle_*$ of (16) is

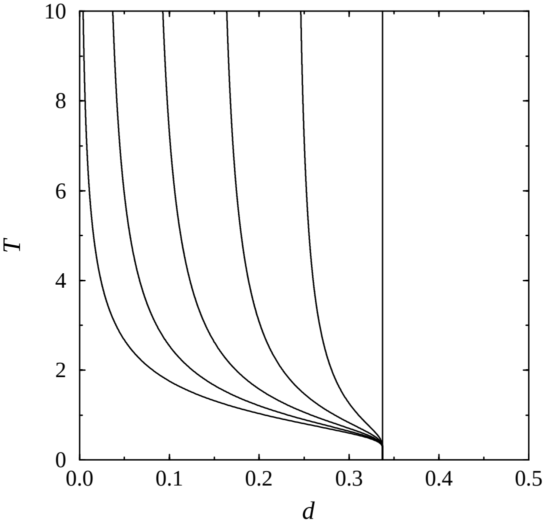

Fig. 5. Phase diagram for mixed multiplicative noise levels: fraction ε at temperature T, fraction (1-ε) at temperature 0; $\varepsilon = (0, 0.2, 0.4, 0.6, 0.8, 1)$ from left to right. For each ε the line separates non-ergodic (left) from ergodic (right) regions. From Ref. [16].

replaced by the average of the right-hand side over the temperature distribution $W(T) = \lim_{N\to\infty} N^{-1} \sum \delta(T - T_i)$. Fig. 5 shows resultant phase diagrams for ensembles with a fraction of agents acting deterministically and the rest at a non-zero temperature.

DS wishes to thank A. Cavagna, J.P. Garrahan, I. Giardina and E. Moro for collaboration which resulted in the joint work reported in the first part of this paper and also for fruitful discussions.

References

[1] D. Challet, Y.-C. Zhang, Physica A 246 (1997) 407.
[2] A. Cavagna, Phys. Rev. E 59 (1998) R3783.
[3] A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, Phys. Rev. Lett. 83 (1999) 4429.
[4] D. Challet, Y.-C. Zhang, Physica A 256 (1998) 514.
[5] R. Savit, R. Manuca, R. Riolo, Phys. Rev. Lett. 82 (1999) 2203.
[6] J.P. Garrahan, E. Moro, D. Sherrington, Phys. Rev. E 62 (2000) R9.
[7] N.F. Johnson, M. Hart, P.M. Hui, Physica A 269 (1999) 1.
[8] D. Challet, M. Marsili, Phys. Rev. E 60 (1999) R6271.
[9] D. Challet, M. Marsili, R. Zecchina, Phys. Rev. Lett. 85 (2000) 5008.
[10] D. Challet, M. Marsili, R. Zecchina, Phys. Rev. Lett. 84 (2000) 1824.
[11] A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, Phys. Rev. Lett. 85 (2000) 5009.
[12] J.P. Garrahan, private communication.
[13] C. de Dominicis, Phys. Rev. B 18 (1978) 4913.
[14] J.A.F. Heimel, A.C.C. Coolen, Phys. Rev. E 63 (2001) 56 121.
[15] A.C.C. Coolen, J.A.F. Heimel, J. Phys. A 34 (2001) 10 783.
[16] A.C.C. Coolen, J.A.F. Heimel, D. Sherrington, Phys. Rev. E 65 (2001) 16 126.