Broken ergodicity and memory in the minority game

2001 J. Phys. A: Math. Gen. 34 L539

(http://iopscience.iop.org/0305-4470/34/40/103)

J. Phys. A: Math. Gen. **34** (2001) L539–L545 PII: S0305-4470(01)27763-4

LETTER TO THE EDITOR

Broken ergodicity and memory in the minority game

J A F Heimel[1] and A De Martino[2]

[1] Department of Mathematics, King's College London, Strand, London WC2R 2LS, UK
[2] International School for Advanced Studies (SISSA/ISAS) and INFM, via Beirut 2-4, 34014 Trieste, Italy

E-mail: heimel@mth.kcl.ac.uk, andemar@sissa.it

Received 9 August 2001
Published 28 September 2001
Online at stacks.iop.org/JPhysA/34/L539

Abstract
We study the dynamics of the 'batch' minority game with market-impact correction using generating functional techniques to carry out the quenched disorder average. We find that the assumption of weak long-term memory, which one usually makes in order to calculate ergodic stationary states, breaks down when the persistent autocorrelation becomes larger than $c_c \simeq 0.772$. We show that this condition, remarkably, coincides with the AT line found in an earlier static calculation. This result suggests a new scenario for ergodicity breaking in disordered systems.

PACS numbers: 02.50.Le, 75.10.Nr

The minority game (MG) models a market of speculators interacting through a simple supply-and-demand mechanism [1, 2]. One of the key behavioural assumptions of the original model is that agents act as so-called price-takers, meaning that at every stage of the game each of them only perceives the aggregate action of all agents, i.e. the total bid. Recently, in [3], a generalization has been introduced in which agents are able to estimate their own contribution to the total bid and use this additional information to adjust their learning dynamics and optimize their performance. The statics of this model has been tackled by spin-glass techniques in [3, 4] along the lines of [5]. The system was found to approximately minimize a disordered Hamiltonian H whose minima could be calculated with the replica method. It was shown that replica-symmetry breaking (RSB) can occur, implying the existence of multiple stationary states.

In this Letter we adapt the dynamical method used in [6] to analyse the 'batch' version of this model. Assuming time-translation invariance, finite integrated response and weak long-term memory [7] we obtain exact results for the stationary state which are in excellent agreement with computer experiments and with earlier static approaches. Moreover, we derive a condition for the continuous onset of memory, where the assumption of weak long-term memory is found to fail while time-translation invariance still holds. This appears to be different from the usual aging scenario in non-ergodic disordered systems. Remarkably, the memory-onset condition coincides with the AT line found in statics.

We begin by recalling the definition of the model. We consider N agents labelled by roman indices. At each iteration round n all agents receive the same information pattern $\mu(n)$ drawn at random with uniform probability from $\{1, \ldots, \alpha N\}$. Each agent has at his disposal S different strategies (labelled by $g = 1, \ldots, S$) to convert the acquired information into a trading decision. Strategies are denoted by αN-dimensional vectors: $\boldsymbol{a}_{ig} = \{a_{ig}^{\mu}\}_{\mu=1}^{\alpha N} \in \{-1, 1\}^{\alpha N}$, where a_{ig}^{μ} is the trading action (e.g. +1 for 'buy', -1 for 'sell') prescribed to agent i by his g-th strategy given receipt of information μ. By assumption, each component a_{ig}^{μ} is selected randomly and independently from $\{-1, 1\}$ with uniform probabilities before the start of the game, for all i, g and μ. This introduces quenched disorder into the model. Each strategy of every agent is given an initial valuation $p_{ig}(0)$, which is updated at the end of every round. At the start of round n, given $\mu(n)$, every agent selects the strategy with the highest valuation, $\widetilde{g}_i(n) = \arg\max p_{ig}(n)$, and subsequently makes a bid according to the trading decision set by the selected strategy: $b_i(n) = a_{i\widetilde{g}_i(n)}^{\mu(n)}$. The total bid at round n is defined as $A(n) = N^{-1/2}\sum_{i=1}^{N} b_i(n)$. Finally, for all i and g all payoffs are updated according to a reinforcement learning dynamics of the form

$$p_{ig}(n+1) = p_{ig}(n) - a_{ig}^{\mu(n)}\left[A(n) - \frac{\eta}{\sqrt{N}}\left(a_{i\widetilde{g}_i(n)}^{\mu(n)} - a_{ig}^{\mu(n)}\right)\right] \tag{1}$$

and agents move to the next round. The first term in square brackets embodies the minority rule, in that the valuation of a strategy is increased every time it predicts the correct minority action, independently of it having been actually used. The term proportional to η adjusts the total bid for the possibility that agent i is not using strategy $\widetilde{g}_i(n)$. For $\eta = 0$ one returns to the original MG, while for $\eta = 1$ the total bid is completely adjusted.

We focus on the case $g = 1, 2$. Introducing the variables $y_i(n) = [p_{i1}(n) - p_{i2}(n)]/2$, as well as the αN-dimensional vectors $\boldsymbol{\omega}_i = (\boldsymbol{a}_{i1} + \boldsymbol{a}_{i2})/2$, $\boldsymbol{\Omega} = N^{-1/2}\sum_{i=1}^{N}\boldsymbol{\omega}_i$ and $\boldsymbol{\xi}_i = (\boldsymbol{a}_{i1} - \boldsymbol{a}_{i2})/2$, and defining $s_i(n) = \mathrm{sgn}[y_i(n)]$ one has

$$y_i(n+1) = y_i(n) - \xi_i^{\mu(n)}[\Omega^{\mu(n)} + \frac{1}{\sqrt{N}}\sum_{j=1}^{N}\xi_j^{\mu(n)}s_j(n) - \frac{\eta}{\sqrt{N}}s_i(n)\xi_i^{\mu(n)}]. \tag{2}$$

Following [6] we study in this paper a 'batch' version of the model, which is obtained by averaging (2) over information patterns:

$$y_i(t+1) = y_i(t) - h_i - \sum_{j=1}^{N} J_{ij}s_j(t) + \eta\alpha s_i(t) + \theta_i(t) \tag{3}$$

where t is a re-scaled time, $h_i = (2/\sqrt{N})\,\boldsymbol{\Omega}\cdot\boldsymbol{\xi}_i$ and $J_{ij} = (2/N)\,\boldsymbol{\xi}_i\cdot\boldsymbol{\xi}_j$. The external field $\theta_i(t)$ has been added for later use. In contrast to the more usual 'on-line' model (2), where the y_i's are updated after every iteration step, in the 'batch' case the updates are made on the basis of the average effect of all possible choices of μ. This modified dynamics yields results for the stationary state which are quantitatively very similar to those of the original model [8]. The theoretical advantage of the 'batch' formulation is that it circumvents the difficulty of constructing a proper continuous time limit. The numerical advantage is that one can simulate larger systems for a longer time. Following [6], one derives the effective non-linear single-agent equation

$$y(t+1) = y(t) - \alpha\sum_{t'\leqslant t}(\mathsf{I}+\mathsf{G})^{-1}_{tt'}s(t') + \alpha\eta s(t) + \sqrt{\alpha}z(t) + \theta(t) \tag{4}$$

where $s(t) = \mathrm{sgn}[y(t)]$ and $z(t)$ is a Gaussian noise with zero mean and temporal correlations given by

$$\langle z(t)z(t')\rangle \equiv H_{tt'} = \sum_{ss'}(\mathsf{I}+\mathsf{G})^{-1}_{ts}(\mathsf{E}+\mathsf{C})_{ss'}(\mathsf{I}+\mathsf{G}^T)^{-1}_{s't'} \tag{5}$$

The matrices C and G appearing here are the noise-averaged single-agent correlation and response functions for the process (4), with elements

$$C_{tt'} = \langle s(t)s(t')\rangle \qquad \text{and} \qquad G_{tt'} = \left\langle \frac{\partial s(t)}{\partial \theta(t')} \right\rangle \tag{6}$$

respectively, while I is the identity matrix and E denotes the matrix with all entries equal to one. The link between the Markovian multi-agent system (2) and the non-Markovian single-agent process (4) is established by the fact that, for $N \to \infty$, $C_{tt'}$ and $G_{tt'}$ become identical to the disorder- and agent-averaged correlation and response functions of (2):

$$C_{tt'} = \frac{1}{N}\sum_{i=1}^{N}\left[s_i(t)s_i(t')\right]_{\mathrm{dis}} \qquad \text{and} \qquad G_{tt'} = \frac{1}{N}\sum_{i=1}^{N}\left[\frac{\partial s_i(t)}{\partial \theta_i(t')}\right]_{\mathrm{dis}}. \tag{7}$$

Equations (4)–(6) describe the dynamics of the system exactly in the $N \to \infty$ limit. We now move to the stationary states of (4) upon making the following assumptions:

- Time-translation invariance (TTI) $\lim_{t\to\infty} C_{t+\tau,t} = C(\tau) \quad \lim_{t\to\infty} G_{t+\tau,t} = G(\tau)$.
- Finite integrated response (FIR) $\lim_{t\to\infty} \sum_{t'\leqslant t} G_{tt'} = \chi < \infty$.
- Weak long-term memory (WLTM) $\lim_{t\to\infty} G(t,t') = 0 \quad \forall t'$ finite.

For the re-scaled quantity $\widetilde{y} = \lim_{t\to\infty} y(t)/t$ one finds

$$\widetilde{y} = -\frac{\alpha s}{1+\chi} + \alpha\eta s + \sqrt{\alpha} z + \theta \tag{8}$$

where $s = \lim_{\tau\to\infty} \tau^{-1}\sum_{t<\tau} \mathrm{sgn}[y(t)]$ and $z = \lim_{\tau\to\infty} \tau^{-1}\sum_{t<\tau} z(t)$, while θ is a static field. The variance of the zero-average Gaussian random variable z can be calculated from (5), yielding

$$\langle z^2\rangle = \lim_{\tau,\tau'\to\infty} \sum_{t\leqslant\tau}\sum_{t'\leqslant\tau'} H_{tt'} = \frac{1+c}{(1+\chi)^2} \tag{9}$$

with the persistent correlation $c \equiv \langle s^2\rangle = \lim_{\tau\to\infty} \tau^{-1}\sum_{t<\tau} C(t)$. The effective agent is 'frozen' if $\widetilde{y} \neq 0$, so that $s = \mathrm{sgn}(\widetilde{y})$ and he is always employing the same strategy. Setting $\theta = 0$, this is easily seen to be the case if $|z| > \gamma$ with $\gamma = \sqrt{\alpha}\left[(1+\chi)^{-1} - \eta\right]$, provided $\gamma \geqslant 0$. He is instead fickle when $\widetilde{y} = 0$ or $|z| < \gamma$, and in this case $s = z/\gamma$. A self-consistent equation for c can now be derived by separating the contribution of the frozen agents from that of the fickle ones. Upon defining $\lambda = \gamma/\sqrt{\langle z^2\rangle}$ one finds

$$c = \langle\Theta(|z|-\gamma)\rangle + \left\langle \Theta(\gamma - |z|)\frac{z^2}{\gamma^2}\right\rangle = \phi + \frac{1}{\lambda^2}\left[\overline{\phi} - \lambda\sqrt{\frac{2}{\pi}}\mathrm{e}^{-\lambda^2/2}\right] \tag{10}$$

where Θ is the step function, $\overline{\phi} = \mathrm{erf}(\lambda/\sqrt{2})$ is the fraction of fickle agents, and $\phi = 1 - \overline{\phi}$ is the fraction of frozen agents. For $\chi = \left\langle\frac{\partial s}{\partial\theta}\right\rangle = \alpha^{-1/2}\left\langle\frac{\partial s}{\partial z}\right\rangle$ one obtains

$$\chi = \frac{1}{\sqrt{\alpha}}\langle\Theta(|z|-\gamma)2\delta(\sqrt{\alpha}z)\rangle + \frac{1}{\gamma\sqrt{\alpha}}\langle\Theta(\gamma-|z|)\rangle = \frac{\overline{\phi}}{\gamma\sqrt{\alpha}}. \tag{11}$$

Equations (9)–(11) form a closed set from which one can solve for ϕ, c and χ for any α and η. Results for c are shown in figure 1.

For negative η, one observes an excellent agreement between theory and experiment for all values of α, implying that none of our assumptions is ever violated.

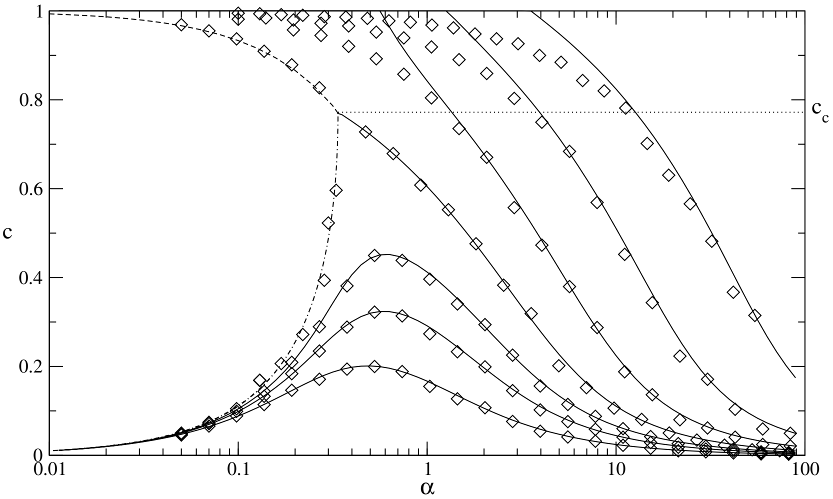

Figure 1. The persistent correlation c as a function of α for different values of η. Lines represent theoretical predictions. Solid lines: from bottom to top, $\eta = -1, -0.5, -0.25, 0, 0.25, 0.5, 0.7$. Dashed line: $y(0) \gg 0, \eta = 0$. Dot-dashed line: $\eta = 0^-$. Diamonds correspond to computer simulations with $\alpha N^2 = 10\,000$, run for 500 time steps and averaged over 50 disorder samples.

When $\eta = 0$, we recover the results of [3,6], which match the simulations perfectly for α larger than the critical value $\alpha_c \simeq 0.3374$. At this point the integrated response χ diverges (FIR is violated) and a transition to a highly non-ergodic regime takes place, where the stationary state depends on the initial conditions $y(0)$. Starting with $y(0) \simeq 0$ leads to a high volatility state, while starting with $|y(0)| \gg 1$ leads to relatively low volatility. The latter regime can be solved using the assumption that χ remains very large for all $\alpha < \alpha_c$. In fact, if $\chi \gg 1$ then $\gamma \simeq \sqrt{\alpha}/\chi$ so that $\overline{\phi} = \alpha$, which is equivalent to $\mathrm{erf}(\lambda/\sqrt{2}) = \alpha$. Solving this for λ and inserting the resulting value in (10), we obtain the top left branch of the $\eta = 0$ curve in figure 1, which is again in excellent agreement with numerical results.

For positive η, one sees that when $c > c_c \simeq 0.77$ our theoretical predictions deviate from the experimental observations, whereas the agreement is perfect for $c < c_c$. Finding no violation of FIR, we have to conclude that either TTI or WLTM is violated. However, we have found no evidence of aging. Therefore we expect the deviations to be related to the breakdown of WLTM only. To find the onset of memory, we split $G_{tt'}$ in its TTI part and its non-TTI part:

$$\lim_{t\to\infty} G_{tt'} = \widetilde{G}(t-t') + \widehat{G}(t,t'). \tag{12}$$

During the initial stages of the game, small perturbations can cause some agents, which would otherwise have remained fickle, to freeze and vice versa, thus creating a persistent part $\widehat{\mathsf{G}}$ in the response function. As the agents freeze, their state (and consequently their contribution to G) becomes independent of t, so that we expect $\lim_{t\to\infty} \widehat{G}(t,t') = \widehat{G}(t')$. After an initial equilibration period, for all frozen agents the difference between the strategy valuations have become very large, so they are virtually insensitive to perturbations. Hence we must assume that $\lim_{t'\to\infty} \widehat{G}(t') = 0$. The fickle agents, however, remain sensitive to small perturbations. The effects will wear out over time (finite response) and are given by $\widetilde{\mathsf{G}}$.

Assuming $\widehat{\mathsf{G}}$ is small, we expand $(\mathsf{I}+\mathsf{G})^{-1}$ in powers of $\widehat{\mathsf{G}}$ up to first order:

$$(\mathsf{I}+\mathsf{G})^{-1} = (\mathsf{I}+\widetilde{\mathsf{G}})^{-1} - \sum_{n=0}^{\infty}\sum_{m=0}^{n-1}(-\widetilde{\mathsf{G}})^m\widehat{\mathsf{G}}(-\widetilde{\mathsf{G}})^{n-m-1} + \mathrm{O}(\widehat{\mathsf{G}}^2). \tag{13}$$

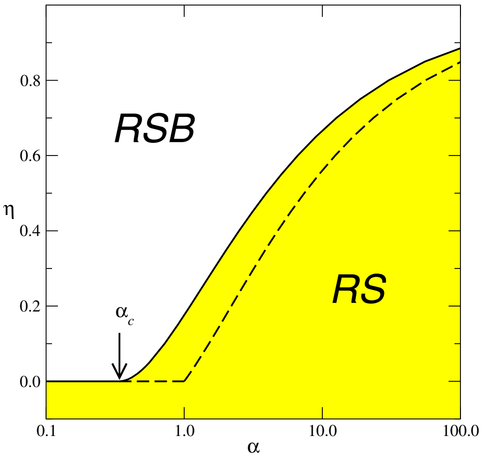

Figure 2. The solid line represents the MO(=AT) line. The dashed line corresponds to the AT line reported in [4].

Defining $\widetilde{\chi} = \sum_t \widetilde{G}(t)$ and $\widehat{\chi} = \sum_t \widehat{G}(t)$, one then finds asymptotically

$$\widetilde{y} = -\alpha \left(\frac{1}{1+\widetilde{\chi}} - \eta \right) s + \sqrt{\alpha} z.$$

$$+\alpha \sum_{n=0}^{\infty} \sum_{m=0}^{n-1} (-\widetilde{\chi})^m \sum_{t'} \widehat{G}(t') \sum_{t''} \left[(-\widetilde{\mathsf{G}})^{n-m-1} \right] (t', t'') s(t'') \tag{14}$$

Using the rectified linear function $f(x) = x$ for $|x| \leqslant 1$ and $\mathrm{sgn}(x)$ otherwise, we see that if $1/(1+\widetilde{\chi}) > \eta$ then

$$s = f\left(\frac{1}{\widetilde{\gamma}} \left[z + \sqrt{\alpha} \sum_{n=0}^{\infty} \sum_{m=0}^{n-1} (-\widetilde{\chi})^m \sum_{t'} \widehat{G}(t') \sum_{t''} \left[(-\widetilde{\mathsf{G}})^{n-m-1} \right] (t', t'') s(t'') \right] \right) \tag{15}$$

$$\widetilde{\gamma} = \sqrt{\alpha} \left(\frac{1}{1+\widetilde{\chi}} - \eta \right).$$

As before, we have $\widetilde{\chi} = \alpha^{-1/2} \left\langle \frac{\partial s}{\partial z} \right\rangle$, whereas

$$\widehat{G}(t) = \left\langle \frac{\partial s}{\partial \theta(t)} \right\rangle = \frac{\sqrt{\alpha}}{\widetilde{\gamma}} \sum_{n=0}^{\infty} \sum_{m=0}^{n-1} (-\widetilde{\chi})^m \sum_{t'} \widehat{G}(t') \sum_{t''} \left[(-\widetilde{\mathsf{G}})^{n-m-1} \right] (t', t'') \widetilde{G}(t'', t). \tag{16}$$

Up to first order in $\widehat{\mathsf{G}}$ one finds $\widehat{\chi} = \widehat{\chi}\widetilde{\chi}\sqrt{\alpha}/[\widetilde{\gamma}(1+\widetilde{\chi})^2] + \mathrm{O}(\widehat{G}^2)$. Although $\widehat{\chi} = 0$ is always a solution of this equation, a bifurcation occurs when $\widetilde{\chi}\sqrt{\alpha}/[\widetilde{\gamma}(1+\widetilde{\chi})^2] = 1$, which is equivalent to $\overline{\phi} = \alpha[1 - \eta(1+\chi)]^2$, and can be written in terms of λ as

$$\lambda^2 [1 + c(\lambda)] = \overline{\phi(\lambda)}. \tag{17}$$

We call this line in the (α, η) plane the memory-onset (MO) line, see figure 2. It coincides remarkably with the AT line (see appendix), and implies that the bifurcation occurs at $c_c \simeq 0.7722$ for $\eta > 0$. Above this value, WLTM can be broken, and indeed one sees from figure 1 that numerical results deviate from our theoretical predictions for $c > c_c$. To give further evidence of memory, we have analysed the time evolution of two identical copies a and b of the system, starting from slightly different initial conditions. We plotted in figure 3 the distance d of the stationary states, given by $(1/N) \sum_i (s_i^a - s_i^b)^2$, where s_i^m is the long-time average of $\mathrm{sgn}(y_i^m)$ $(m = a, b)$, versus the persistent autocorrelation of copy a, c^a. As c^a approaches c_c, the two copies end up in different stationary states, proving that they remember

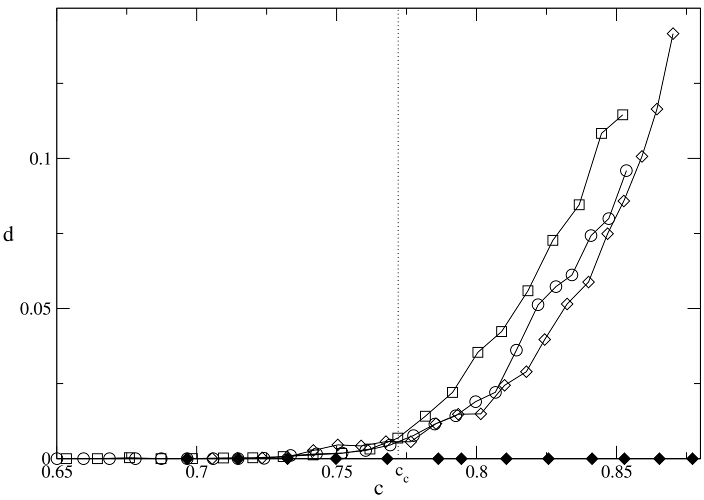

Figure 3. Distance between the stationary states of two identical copies of the system (see text) as a function of their persistent autocorrelation. Simulations are for various levels of η with $N = 450$, averaged over 100 samples. Open markers correspond to a perturbation at $t = 0$ (○, □, ◇ for $\alpha = 1, 2, 4$, respectively). Closed markers correspond to a perturbation at $t = 500$ (◆ for $\alpha = 2$). All simulations are run up to 500 steps after the perturbation occurs. Time averages are over the last 300 steps.

initial conditions[3]. At the same time, if a perturbation is applied much later during the run the copies end up in the same stationary state, indicating that indeed $\widehat{G}(t') \to 0$ as $t' \to \infty$.

Summarizing, we have shown that in this model the usual connection between broken ergodicity and broken TTI (aging), as seen for instance in mean-field spin glasses [9], does not occur. In contrast, we derive from the dynamics a condition for breakdown of WLTM and continuous onset of memory within the TTI regime, which is found to be equivalent to the AT line found in the static approach. This remarkable deviation from the well-known RSB/aging picture is possibly due to the fact that the microscopic dynamics of our model does not satisfy detailed balance.

We gratefully acknowledge support from and useful discussions with A C C Coolen and M Marsili, and with S Franz. We also thank SISSA and King's College London for reciprocal hospitality. This work originated at the International Seminar on Statistical Mechanics of Information Processing in Cooperative Systems (Dresden, March 2001).

Appendix

While investigating the possible relation between our MO line and replica-symmetry breaking, it became apparent that in the very final step of the AT line calculation in [4] a small error has occurred. It was found that the replica-symmetric solution becomes unstable when

$$\lim_{\beta\to\infty} \left\langle \beta^2 \left(\langle s^2 \rangle - \langle s \rangle^2 \right)^2 \right\rangle_z = (1+\chi)^2 \tag{18}$$

[3] The slight bump that occurs before c_c is likely due to the fact that in our simulation the perturbation can not be infinitesimal, but is at least $1/N$.

where $\langle f(s)\rangle = Z_\beta(z)^{-1}\int_{-1}^{1} f(s)e^{-\beta V_z(s)}\mathrm{d}s$ and $Z_\beta(z) = \int_{-1}^{1} e^{-\beta V_z(s)}\mathrm{d}s$, with $V_z(s) = \frac{1}{2}\gamma s^2 - zs$. The brackets $\langle\cdots\rangle_z$ denote a Gaussian average over z having zero mean and variance $\left\langle z^2\right\rangle_z = (1+q)/(1+\chi)^2$, q being the overlap between two different replicas (off-diagonal overlap matrix element). We have absorbed a spurious factor $\sqrt{\alpha}$ in β. If we now define $F(z) = -\lim_{\beta\to\infty}\beta^{-1}\log Z_\beta(z)$, the AT line can be written as $\left\langle F''(z)^2\right\rangle_z = (1+\chi)^2$. By Laplace's method we find $F(z) = V_z(s_0)$, s_0 being the minimum of V_z in $[-1, 1]$. For $|z/\gamma| < 1$, s_0 lies inside this interval and $V_z(s_0) = z^2/(2\gamma)$, while for $|z/\gamma| > 1$ s_0 is on the border and $V_z(s_0) = \gamma/2 - |z|$. This gives second derivatives that are $-\gamma^{-1}$ and 0, respectively. The AT line is therefore given by $\langle\gamma^{-2}\Theta(1-|z/\gamma|)\rangle_z + \langle 0\ \Theta(|z/\gamma|-1)\rangle_z = (1+\chi)^2$. Recognizing the non-vanishing term on the lhs as the fraction $\overline{\phi}$ of fickle agents, we find

$$\alpha[1-\eta(1+\chi)]^2 = \overline{\phi} \tag{19}$$

similar to the result of [4] where in place of $\overline{\phi}$ a 1 was reported. Written in terms of λ, the AT line is identical to the MO line (17). We learned that the AT line can also be derived from the dynamical stability of equation (2) [10].

References

[1] Challet D and Zhang Y-C 1997 *Physica* A **246** 407

[2] See the web page There is an extensive and commented overview of the existing literature at `www.unifr.ch/econophysics/minority`

[3] Marsili M, Challet D and Zecchina R 2000 *Physica* A **280** 522

[4] De Martino A and Marsili M 2001 *J. Phys. A: Math. Gen.* **34** 2525

[5] Challet D, Marsili M and Zecchina R 2000 *Phys. Rev. Lett.* **84** 1824

[6] Heimel J A F and Coolen A C C 2001 *Phys. Rev.* E **63** 56121

[7] Cugliandolo L F and Kurchan J 1995 *Phil. Mag.* B **71** 501

[8] Coolen A C C and Heimel J A F 2001 in preparation

[9] Bouchaud J-P, Cugliandolo L F, Kurchan J and Mézard M 1998 *Spin Glasses and Random Fields* ed A P Young (Singapore: World Scientific) pp 161–224

[10] Marsili M 2001 private communication