

Non-equilibrium statistical mechanics of Minority Games

A C C Coolen

Department of Mathematics, King's College London
The Strand, London WC2R 2LS, UK

Abstract. In this paper I give a brief introduction to a family of simple but non-trivial models designed to increase our understanding of collective processes in markets, the so-called Minority Games, and their non-equilibrium statistical mathematical analysis. Since the most commonly studied members of this family define disordered stochastic processes without detailed balance, the canonical technique for finding exact solutions is found to be generating functional analysis a la De Dominicis, as originally developed in the spin-glass community.

Keywords: non-equilibrium statistical mechanics, disordered systems

AMS subject classification: 60K35

1. Introduction

The Minority Game (MG) [1], a variation on the so-called El-Farol bar problem [2], was designed to understand the cooperative phenomena observed in markets. It describes agents who each make a binary decision at every point in time, e.g. whether to buy or sell. Profit is made by those who find themselves in the minority group, i.e. who end up buying when most wish to sell, or vice versa. The agents take their decisions on the basis of an arsenal of individual 'strategies', from which each agent aims to select the optimal one (i.e. the strategy that leads to the largest number of minority decisions). The dynamic equations of the MG describe the stochastic evolution of these strategy selections. Each agent wishes to make profit, but the net effect of his/her trading actions is defined fully in terms of (or relative to) the actions taken by the other agents. Hence the model has a significant amount of built-in frustration. Although in its simplest form (as discussed in this paper) the MG is a crude simplification of real markets, there is now also a growing number of more realistic MG spin-off models [3].

For the non-equilibrium statistical mechanicist the MG is an intriguing disordered system. It has been found to exhibit non-trivial behaviour, e.g. dynamical phase transitions separating an ergodic from a highly non-ergodic regime, and to pose a considerable mathematical challenge. Its stochastic equations do not obey detailed balance, so there is no equilibrium state and understanding the model requires solving its dynamics. In this paper I first give an introduction to the properties and the early

studies of the MG, seen through the eyes of a non-equilibrium statistical mechanistic. The second part of the paper is a review of the more recent applications to the MG of the generating functional analysis methods of De Dominicis [4]. These have led to exact results in the form of closed macroscopic order parameter equations, and to clarification of a number of unresolved issues and debates regarding the nature of the microscopic laws and various previously proposed approximations.

2. Definitions and properties

We consider a community of N agents, labeled $i = 1 \dots N$ (the ‘market’). At each discrete iteration step ℓ each agent i must take a binary decision $s_i(\ell) \in \{-1, 1\}$ (e.g. ‘sell’ or ‘buy’). The agents make these decisions on the basis of public information $I(\ell)$ which they are given at stage ℓ (state of the market, political developments, weather forecasts, etc), which is taken from some finite discrete set $\Omega = \{I_1, \dots, I_p\}$ of size $|\Omega| = p$. For future convenience we will also define the relative size $\alpha = p/N$. Note that all agents always receive the same information. Profit is made at iteration step ℓ by those agents who find themselves in the minority group (by those who wish to buy when most aim to sell, or vice versa), i.e. by those i for which $s_i(\ell)[\sum_j s_j(\ell)] < 0$.

The actual conversion of the observation of $I(\ell)$ into the N trading decisions $\{s_1(\ell), \dots, s_N(\ell)\}$ is defined via so-called trading ‘strategies’, which are basically look-up tables. Every agent i has S decision making strategies $\mathbf{R}_{ia} = (R_{ia}^1, \dots, R_{ia}^p) \in \{-1, 1\}^p$, with $a \in \{1, \dots, S\}$. If strategy a is the one used by agent i at iteration step ℓ , and the public information $I(\ell)$ is found to be $I(\ell) = I_{\mu(\ell)}$, then agent i will locate the appropriate entry in the table \mathbf{R}_{ia} and take the decision

$$s_i(\ell) = R_{ia}^{\mu(\ell)} \quad (1)$$

Thus, given knowledge of the active strategies of all N agents, the responses of all agents to the arrival of public information are fully deterministic.

What remains in defining the model is giving a recipe for deciding the choice of decision strategy \mathbf{R}_{ia} from the arsenal of S possibilities as a function of time, for all agents. This is done as follows. The agents aim to select or discover profitable strategies within their arsenals, by keeping track of the strategies’ performance in the market. A strategy \mathbf{R}_{ia} is profitable for agent i at stage ℓ of the game (i.e. it puts agent i in the minority group) if and only if $R_{ia}^{\mu(\ell)} = -\text{sgn}[\sum_j s_j(\ell)]$. Hence we can measure for each agent i the cumulative performance of strategy \mathbf{R}_{ia} by a quantity p_{ia} , which is updated at every stage ℓ of the game according to

$$p_{ia}(\ell + 1) = p_{ia}(\ell) - \frac{\eta}{N} R_{ia}^{\mu(\ell)} \sum_j s_j(\ell) \quad (2)$$

with $\eta > 0$ (the scaling factor N is introduced for future mathematical convenience). Each agent i now chooses at every stage ℓ of the game the strategy $a_{i\ell}$ which has the best cumulative performance $p_{ia}(\ell)$ at that point in time.

In summary, upon insertion of (1) into (2) and given the external information stream




Figure 1. Dependence of the volatility σ on the re-scaled size $\alpha = p/N$ of the external information set Ω . The markers refer to numerical simulations following different initial conditions (whose details will be explained later), with $S = 2$ (two strategies per agent) and $N = 4000$. The value $\sigma = 1$ would have been obtained for purely random decision making.

defined by $\{\mu(\ell)\}$, the dynamics of the MG is defined fully by the following two coupled non-linear equations for the evolving strategy performance measures $\{p_{ia}\}$ (with $A(\ell)$ giving the total market bid at stage ℓ):

$$p_{ia}(\ell + 1) = p_{ia}(\ell) - \frac{\eta}{N} R_{ia}^{\mu(\ell)} A(\ell) \quad A(\ell) = \sum_j R_{ja_j}^{\mu(\ell)} \quad (3)$$

$$a_i(\ell) = \arg \max_{a \in \{1, \dots, S\}} p_{ia}(\ell) \quad (4)$$

In the original model [1] the information $\mu(\ell)$ was defined as a deterministic function of the macroscopic market history, i.e. of the terms $A(\ell') = \sum_j s_j(\ell') = \sum_j R_{ja_j}^{\mu(\ell')}$ with $\ell' < \ell$: $\mu(\ell) = K[A(\ell - 1), A(\ell - 2), \dots]$, for some function $K[\dots]$. In this formulation the process (3,4) is fully deterministic but non-local in time. Furthermore, there is quenched disorder (the realization of the look-up tables $\{\mathbf{R}_j\}$) and a large degree of built-in competition and frustration.

The main quantity of interest to those caring about real markets is the so-called volatility, the measure of macroscopic market fluctuations. In the present model the appropriate object is the re-scaled and time-averaged variance σ^2 of the total bid $A(\ell)$ (since the time-averaged expectation value of the total bid is zero):

$$\sigma^2 = \lim_{\tau \rightarrow \infty} \frac{1}{\tau N} \sum_{\ell=1}^{\tau} A^2(\ell) \quad (5)$$

This quantity was found in simulations to depend in a nontrivial way on the relative size α of the set Ω , similar to e.g. Figure 1, which seemed to suggest to agents the possibility of market prediction (since random decision making corresponds to $\sigma = 1$)

and to theorists the presence of a non-equilibrium phase transition. Since simulations also indicated that the behaviour of the volatility did not depend qualitatively on the value of S (the number of strategies per agent, provided this number remains finite as $N \rightarrow \infty$), most of the research has concentrated on the simplest nontrivial case $S = 2$. Here it is possible to reduce the dimensionality of the problem further upon defining

$$q_i = \frac{p_{i1} - p_{i2}}{2}, \quad \xi_i^\mu = \frac{R_{i1}^\mu - R_{i2}^\mu}{2}, \quad \Omega_\mu = \frac{\sum_j [R_{j1}^\mu + R_{j2}^\mu]}{2\sqrt{N}} \quad (6)$$

We can now write the process (3,4) in the deceptively compact form

$$q_i(\ell + 1) = q_i(\ell) - \frac{\eta}{\sqrt{N}} \xi_i^{\mu(\ell)} \left[\Omega_{\mu(\ell)} + \frac{1}{\sqrt{N}} \sum_j \xi_j^{\mu(\ell)} \text{sgn}[q_j(\ell)] \right] \quad (7)$$

which in the standard MG is non-local in time due to the dependence of $\mu(\ell) = K[A(\ell-1), A(\ell-2), \dots]$ (with $A(\ell') = \Omega_{\mu(\ell')} + \frac{1}{\sqrt{N}} \sum_j \xi_j^{\mu(\ell')} \text{sgn}[q_j(\ell')]$) on the past of the system.

3. Early studies and generalizations

The first major step forward after the MG's introduction in [1] was [5], where it was shown that the volatility σ remains largely unchanged if, instead of being calculated from the true market history, the variables $\mu(\ell)$ are simply drawn at random from $\{1, \dots, \alpha N\}$ (at every iteration step ℓ , independently and with equal probabilities). Apparently the agents are not predicting the market: the observed phenomena are the result of complex collective processes, which depend only on all agents responding to the *same* information (whether correct or nonsensical). The mathematical consequences of having the numbers $\mu(\ell)$ drawn at random, however, are profound: the equations (7) now describe a disordered Markov process (without detailed balance) and analysis comes within reach. In [6] it was shown that the relevant quantities studied in the MG can be scaled such that they become independent of the number of agents N when this number becomes very large. An interesting generalization of the game was the introduction of agents' decision noise [7], which was shown not only to improve worse than random behaviour but also, more surprisingly, to be able to make it better than random[‡]. In [10] one finds the first observation in the MG of the so-called 'frozen agents'. These are agents for which asymptotically $q_i(\ell) \sim \ell$ ($\ell \rightarrow \infty$), in the language of (7), and who in view of the definition $q_i = \frac{1}{2}[p_{i1} - p_{i2}]$ will eventually end up always playing the same strategy. The four papers [5, 6, 7, 10] more or less paved the way for statistical mechanical theory.

The first attempt at solution of the MG is found in [11, 12], whose authors studied the MG with decision noise and argued that for $N \rightarrow \infty$, and after appropriate temporal course-graining, one is allowed to neglect the fluctuations in $\text{sgn}[q_i(\ell) + z_i(\ell)]$ and concentrate on the expectation values $m_i(\ell) = \int dz P(z) \text{sgn}[q_i(\ell) + z]$. In terms

[‡] Using a phenomenological theory for the volatility, based on so-called 'crowd-anticrowd' cancellations this effect was partially explained in [8, 9].

of the re-scaled time $t = \ell/N$ the authors then derived deterministic equations in the limit $N \rightarrow \infty$, of the form $dm_i(t)/dt = f_i[\mathbf{m}(t)]$, which were shown to have a Lyapunov function. This allowed for a ground state analysis using equilibrium statistical mechanics, in combination with standard replica theory [13] to carry out the disorder average (based on the introduction of a fictitious temperature which is sent to zero *after* having taken the limits $N \rightarrow \infty$ and $n \rightarrow 0$, where n denotes the replica dimension, similar to e.g. [14]). The authors of [11, 12] found a phase transition at $\alpha \downarrow \alpha_c \approx 0.33740$ (in agreement with simulation evidence, see e.g. Figure 1), but ran into mathematical difficulties for $\alpha < \alpha_c$.

The authors of [11, 12] then initiated a discussion [15, 16] about the validity of the assumptions in [11, 12] and the reliability of the simulations in [7]. Following a further study [17], in which a Fokker-Planck equation was derived for the MG (and solved numerically), and where fluctuations were claimed not to be negligible, this discussion appeared to be put to rest in [18] with the conclusions that the simulations of [7] had not yet equilibrated (acknowledged already in [16]), and that the microscopic fluctuations in the MG can not be neglected after all (except perhaps in special cases). It should be noted, however, that the microscopic definitions of [7, 16, 17] are different from those of [11, 15, 18]. This state of affairs was unsatisfactory: there was still no exact solution of the MG (whether for statics or dynamics), the status of the assumptions and approximations made by the various authors remained unclear, in both [17] and [18] Fokker-Planck equations are derived for the MG, but with different diffusion matrices, and (finally) there were conflicting statements on the differences between multiplicative and additive decision noise with regard to the volatility. Clearly, since the MG defines a disordered stochastic process without detailed balance, the natural analytical approach is to study its dynamics. In the remainder of this paper I show how the generating functional techniques developed by De Dominicis [4] can be used to solve various versions of the MG analytically. Full details of this analysis can be found in the trio [19, 20, 21].

4. Solution of the Batch MG

In order to circumvent the debates relating to the temporal course graining of the MG process and concentrate fully on the disorder, the first generating functional analysis study of the MG [19] was carried out for a modified so-called ‘batch’ version of the model. Instead of the equations (7) (where state changes follow each randomly drawn information index $\mu(\ell)$), in the batch MG the dynamics is defined directly in terms of an average over all possible choices $\mu \in \{1, \dots, \alpha N\}$ for the external information:

$$q_i(\ell + 1) = q_i(\ell) - \frac{1}{\alpha N} \sum_{\mu=1}^{\alpha N} \xi_i^\mu \left\{ \Omega_\mu + \frac{1}{\sqrt{N}} \sum_j \xi_j^\mu \operatorname{sgn}[q_j(\ell)] \right\} + \theta_i(\ell) \quad (8)$$

Here we have also chosen $\eta = \sqrt{N}$ (to ensure $\mathcal{O}(N^0)$ characteristic time-scales for the batch version of the MG). This defines a disordered deterministic map, since the

stochasticity has now been removed. An external force $\theta_i(\ell)$ was added in order to define response functions later.

Generating functional analysis a la De Dominicis [4] is based on evaluation of the following disorder-averaged§ generating functional:

$$\overline{Z[\boldsymbol{\psi}]} = \overline{\left\langle e^{-i \sum_{i=1}^N \sum_{t \geq 0} \psi_i(t) q_i(t)} \right\rangle} \quad (9)$$

From this object one can obtain the main disorder-averaged quantities of interest, such as individual averages, multiple-time covariances, and response functions, e.g.

$$\overline{\langle q_i(\ell) \rangle} = \lim_{\boldsymbol{\psi} \rightarrow \mathbf{0}} \frac{i \partial \overline{Z[\boldsymbol{\psi}]} }{\partial \psi_i(\ell)} \quad (10)$$

$$\overline{\langle q_i(\ell) q_j(\ell') \rangle} = - \lim_{\boldsymbol{\psi} \rightarrow \mathbf{0}} \frac{\partial^2 \overline{Z[\boldsymbol{\psi}]} }{\partial \psi_i(\ell) \partial \psi_j(\ell')} \quad (11)$$

$$\frac{\partial \overline{\langle q_i(\ell) \rangle}}{\partial \theta_j(\ell')} = \lim_{\boldsymbol{\psi} \rightarrow \mathbf{0}} \frac{i \partial^2 \overline{Z[\boldsymbol{\psi}]} }{\partial \psi_i(\ell) \partial \theta_j(\ell')} \quad (12)$$

One next defines suitable (time-dependent) fields which must depend linearly on an extensive number of the statistically independent disorder variables $\{R_{ia}^\mu\}$, and which drive the dynamics (in spin-glass models these would have been the actual local magnetic fields), such as $x^\mu(\ell) = \sqrt{2/N} \sum_i \text{sgn}[q_i(\ell)] \xi_i^\mu$. One then writes (9) as a sum (or integral) over all possible paths taken by these fields. In practice this is done via the insertion into (9) of integrals over appropriate δ -distributions, e.g. of

$$1 = \prod_{\mu=1}^{\alpha N} \prod_{\ell \geq 0} \left\{ \int dx^\mu(\ell) \delta \left[x^\mu(\ell) - \frac{\sqrt{2}}{\sqrt{N}} \sum_i \text{sgn}[q_i(\ell)] \xi_i^\mu \right] \right\} \quad (13)$$

This procedure is carried out until all disorder variables have been concentrated into such fields; here this requires further fields in addition to the $x^\mu(\ell)$. Upon writing the δ -functions in expressions such as (13) in integral form, the independent disorder variables $\{R_{ia}^\mu\}$ will all occur linearly in exponents, so that the disorder average can be carried out. At the end one finds an expression involving an integral over all possible values of the single-site correlation and response functions

$$C_{tt'} = \frac{1}{N} \sum_{i=1}^N \overline{\langle \text{sgn}[q_i(t)] \text{sgn}[q_i(t')] \rangle} \quad (14)$$

$$G_{tt'} = \frac{1}{N} \sum_{i=1}^N \frac{\partial}{\partial \theta_i(t')} \overline{\langle \text{sgn}[q_i(t)] \rangle} \quad (15)$$

and a further set of related order parameter kernels $\{\hat{C}_{tt'}, \hat{G}_{tt'}, L_{tt'}, \hat{L}_{tt'}\}$:

$$\overline{Z[\boldsymbol{\psi}]} = \int \mathcal{D}C \mathcal{D}\hat{C} \mathcal{D}G \mathcal{D}\hat{G} \mathcal{D}L \mathcal{D}\hat{L} e^{N\Psi[C, \hat{C}, G, \hat{G}, L, \hat{L}]} \quad (16)$$

§ Disorder averaging will be denoted as $\overline{[\dots]}$. The disorder is in the variables ξ_i^μ and Ω_μ , see the definitions (6), with $\mu \in \{1, \dots, \alpha N\}$ and with each of the $2\alpha N^2$ strategy table entries R_{ia}^μ drawn independently at random from $\{-1, 1\}$ with equal probabilities.

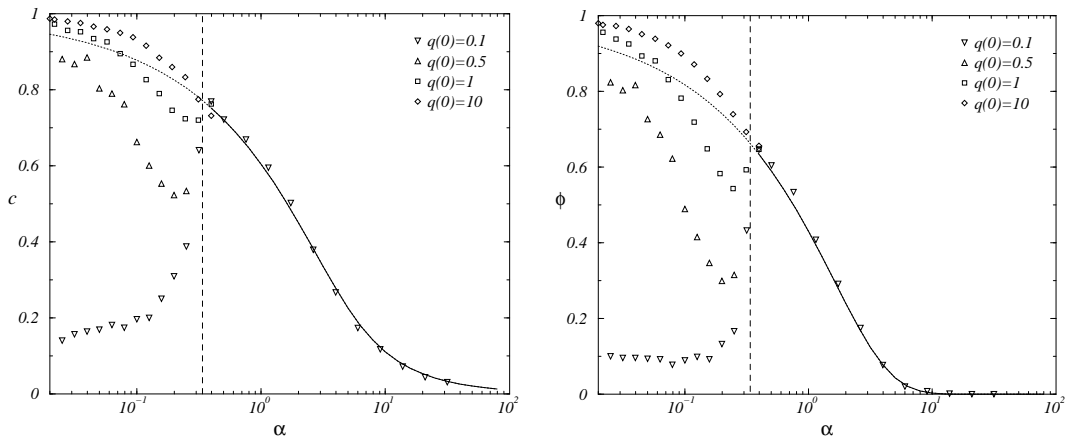


Figure 2. The persistent covariance $c = \lim_{t \rightarrow \infty} C(t)$ (left picture) and the fraction ϕ of frozen agents (right picture) in the stationary state. The markers are obtained from individual simulation runs, with $N = 4000$ and various homogeneous initial conditions (where $q_i(0) = q(0) \forall i$), measured after 1000 iteration steps. The solid curves to the right of the critical point (vertical dashed line in each picture) are the theoretical predictions, given by the solution of (20,21). The dotted curves to the left of the critical point are their continuations into the $\alpha < \alpha_c$ regime (where they should no longer be correct).

with the abbreviation $\mathcal{D}C \equiv \prod_{tt' \geq 0} dC_{tt'}$, and with $\Psi[\dots] = \mathcal{O}(N^0)$. In the limit $N \rightarrow \infty$ the integral in (16) is evaluated by steepest descent, giving closed dynamical equations for the kernels $\{C, \hat{C}, G, \hat{G}, L, \hat{L}\}$, from which $\{\hat{C}, \hat{G}, L, \hat{L}\}$ can be eliminated. This leaves a theory described by closed equations for C and G only, which turns out to describe an effective stochastic ‘single agent’ process of the form

$$q(t+1) = q(t) + \theta(t) - \alpha \sum_{t' \leq t} [\mathbf{I} + G]_{tt'}^{-1} \text{sgn}[q(t')] + \sqrt{\alpha} \eta(t) \quad (17)$$

In (17) one has a retarded self-interaction and a zero-average Gaussian noise $\eta(t)$ with

$$\langle \eta(t) \eta(t') \rangle = \left[(\mathbf{I} + G)^{-1} D (\mathbf{I} + G^\dagger)^{-1} \right]_{tt'} \quad D_{tt'} = 1 + C_{tt'} \quad (18)$$

The equations from which to solve the two kernels $\{C, G\}$ self-consistently are defined in terms of the non-Markovian single-agent process (17,18) as follows: for all $t, t' \in \mathbb{N}$

$$C_{tt'} = \langle \text{sgn}[q(t)] \text{sgn}[q(t')] \rangle, \quad G_{tt'} = \frac{\partial}{\partial \theta(t')} \langle \text{sgn}[q(t)] \rangle \quad (19)$$

Finding the general solution of (19) analytically is extremely hard; one has to restrict oneself in practice to working out the first few time steps, to stationary solutions, or to numerical solution.

If, for instance, one investigates asymptotic stationary solutions of the form $C_{tt'} = C(t-t')$ and $G_{tt'} = G(t-t')$, and furthermore assumes absence of anomalous response (i.e. a finite value for the integrated response $\chi = \sum_{t>0} G(t)$) it is not difficult




Figure 3. Evidence for the existence of a critical initial strategy valuation $q_c(0)$ below which a high-volatility solution $\sigma \sim \alpha^{-1/2}$ ($\alpha \rightarrow 0$) exists. Connected markers show the results of individual simulations, with $N = 4000$ and homogeneous initial conditions $q_i(0) = q(0) \forall i$, measured after 1000 iteration steps. The data are compatible with the prediction $q_c(0) \approx 0.242$ (dashed line).

[19] to obtain an exact closed equation for the persistent correlation $c = \lim_{t \rightarrow \infty} C(t)$:

$$c = 1 - \left(1 - \frac{1+c}{\alpha}\right) \operatorname{erf} \left[\sqrt{\frac{\alpha}{2(1+c)}} \right] - \sqrt{\frac{2(1+c)}{\pi\alpha}} e^{-\frac{\alpha}{2(1+c)}}. \quad (20)$$

from which c follows as a function of the parameter α . Similarly one can find exact expressions for the fraction ϕ of ‘frozen’ agents and for the integrated response χ :

$$\phi = 1 - \operatorname{erf} \left[\sqrt{\alpha/2(1+c)} \right] \quad (21)$$

$$\chi = (1 - \phi)/(\alpha - 1 + \phi) \quad (22)$$

For $\alpha \rightarrow \infty$ one finds $\lim_{\alpha \rightarrow \infty} c = \lim_{\alpha \rightarrow \infty} \phi = \lim_{\alpha \rightarrow \infty} \chi = 0$. It follows from (22) that the assumption $\chi < \infty$ underlying (20,21,22) breaks down when $\alpha = 1 - \phi$. This defines a critical value α_c signaling the onset of non-ergodicity, which can be calculated in combination with (21), giving $\alpha_c \approx 0.33740$ (the same value obtained earlier in [11, 12]). The results of solving (20) and (21) numerically, as a function of α , are shown in Figure 2. One observes perfect agreement with simulations for $\alpha > \alpha_c$, and, as expected, disagreement and ergodicity breaking for $\alpha < \alpha_c$ (where the condition $\chi < \infty$ is violated). The volatility is found not to be a natural order parameter of the MG (nor expressible directly in terms of the solution of (20,21), not even for $\alpha > \alpha_c$; its calculation required approximations, resulting in Figure 1 (shown together with simulation results of the same specifications as those in Figure 2).

It should be emphasized that for $\alpha < \alpha_c$ the macroscopic equations (17,18,19) are still exact; finding their stationary solutions analytically, however, is much more difficult. Some of the results which have been obtained for the $\alpha < \alpha_c$ regime, for instance, are scaling properties of solutions for $\alpha \rightarrow 0$ (shown in Figure 1). Low volatility solutions $\sigma \sim \alpha^{1/2}$ ($\alpha \rightarrow 0$) are found following initial conditions with large

$|q_i(0)|$, whereas high-volatility solutions with $\sigma \sim \alpha^{-\frac{1}{2}}$ ($\alpha \rightarrow 0$) are found following initial conditions with small $|q_i(0)|$. Upon making some simple approximations one can also calculate an expression for the critical initial value for the $|q_i(0)|$, $q_c(0) \approx 0.242$, bordering the region of existence of the high-volatility solution, see Figure 3.

5. Solution of the Batch MG with decision noise

The Batch MG with decision noise is obtained upon making in (8) the replacement $\text{sgn}[q_i(t)] \rightarrow \sigma[q_i(\ell), z_i(\ell)|T_i]$, with the $\{z_i(\ell)\}$ representing zero-average and unit-variance identically distributed and independent random variables. The parameters $T_i \geq 0$ measure the noise levels of the individual agents. The function $\sigma[\dots|T]$ obeys $\sigma[q, z|T] \in \{-1, 1\}$, $\sigma[q, z|0] = \text{sgn}[q]$ (for $T \rightarrow 0$ we return to the previous model) and $\int dz P(z)\sigma[q, z|\infty] = 0$ ($T \rightarrow \infty$ implies purely random decisions). Examples are

$$\text{additive noise :} \quad \sigma[q, z|T] = \text{sgn}[q + Tz] \quad (23)$$

$$\text{multiplicative noise :} \quad \sigma[q, z|T] = \text{sgn}[q] \text{sgn}[1 + Tz] \quad (24)$$

The deterministic equations (8) are thus replaced by the following stochastic ones:

$$q_i(\ell + 1) = q_i(\ell) - \frac{1}{\alpha N} \sum_{\mu=1}^{\alpha N} \xi_i^\mu \left\{ \Omega_\mu + \frac{1}{\sqrt{N}} \sum_j \xi_j^\mu \sigma[q_j(\ell), z_j(\ell)|T_j] \right\} + \theta_i(\ell) \quad (25)$$

The analytical procedure outlined for the previous model (8) is adapted quite easily to its generalization (25). Again one finds, in the limit $N \rightarrow \infty$, closed equations for the disorder-averaged correlation and response functions (14,15), describing an effective stochastic ‘single agent’ process:

$$q(t + 1) = q(t) + \theta(t) - \alpha \sum_{t' \leq t} [\mathbf{I} + G]_{tt'}^{-1} \sigma[q(t'), z(t')|T] + \sqrt{\alpha} \eta(t) \quad (26)$$

(parametrized by a noise strength T , so that we must denote averages henceforth as $\langle \dots \rangle_T$). The covariances of the zero-average Gaussian noise $\eta(t)$ are still given by (18); the independent random variables $\{z(t)\}$ are distributed exactly as the site-variables $\{z_i(\ell)\}$ in (25). The order parameter equations (19) are now found to be replaced by

$$C_{tt'} = \int_0^\infty dT W(T) \langle \text{sgn}[q(t)] \text{sgn}[q(t')] \rangle_T \quad (27)$$

$$G_{tt'} = \int_0^\infty dT W(T) \frac{\partial}{\partial \theta(t')} \langle \text{sgn}[q(t)] \rangle_T \quad (28)$$

with the distribution of noise strengths

$$W(T) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \delta[T - T_i] \quad (29)$$

Making the choice $W(T) = \delta(T)$ (i.e. deterministic decision making) immediately brings us back to (17,18,19), as it should.




Figure 4. Transition lines for the Batch MG with multiplicative decision noise (24), $P(z) = (2\pi)^{-1/2}e^{-z^2/2}$, and $W(T') = \epsilon\delta[T' - T] + (1 - \epsilon)\delta[T']$ with $\epsilon \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$ (from right to left). Each line is defined by the condition $\chi \equiv \sum_{t>0} G(t) = \infty$ and separates a non-ergodic phase (to the left) from an ergodic phase (to the right). For additive noise our theory predicts the T -independent (vertical) line, i.e. the $\epsilon = 0$ curve, for any choice of $W(T)$.

In the present model one can as before calculate stationary solutions without anomalous response, and calculate order parameters such as c and ϕ as a function of α and the distribution $W(T)$, and for different types of decision noise. A phase transition to a highly non-ergodic state will again be marked by $\chi \equiv \sum_{t>0} G(t) = \infty$. Here I restrict myself to showing phase diagrams for some simple choices, see Figure 4. The theory reveals that, in the ergodic regime, for additive decision noise (23) the location of the $\chi = \infty$ transition is completely independent of $W(T)$; the same is true for the values of the order parameters c and ϕ . Calculating the solution of equations (26,18,27,28) for the first few time-steps [20] also reveals the mechanism of how the introduction of decision noise can indeed have the counter-intuitive effect (observed already in [7]) of reducing the market volatility σ , by damping global oscillations.

6. Solution of the On-line Minority Game

Having carried out the generating functional analysis of the two previous batch versions of the MG (with $S = 2$, and with deterministic or noisy decision making), the obvious next step is to return to the original (more complicated) dynamics (7) where state changes followed individual presentations of external information, but now generalized

|| The explanation for this seems to be that the $\chi = \infty$ transition is determined purely by the properties of the ‘frozen agents’. These agents, for which $q_i(\ell) \sim \ell$ as $\ell \rightarrow \infty$ will asymptotically always overrule the noise term in (23). In contrast, for multiplicative noise (24) the noise will retain the potential to change decisions, even for ‘frozen’ agents.

to include decision noise:

$$q_i(\ell + 1) = q_i(\ell) - \frac{\eta}{\sqrt{N}} \xi_i^{\mu(\ell)} \left\{ \Omega_{\mu(\ell)} + \frac{1}{\sqrt{N}} \sum_j \xi_j^{\mu(\ell)} \sigma[q_j(\ell), z_j(\ell)] \right\} + \theta_i(\ell) \quad (30)$$

The $\mu(\ell)$ are drawn randomly and independently from $\{1, \dots, \alpha N\}$, introducing a second element of stochasticity on top of the decision making noise variables $\{z_j(\ell)\}$. However, the individual (noisy) updates are now small ($\mathcal{O}(N^{-\frac{1}{2}})$), so for $N \rightarrow \infty$ we may attempt a continuous-time description. This switch to continuous time (in particular: whether upon doing so the fluctuations can be taken as Gaussian, or even be discarded) was the subject of the discussion in [15, 16] and subsequent papers.

The temporal regularization problem can be resolved using a procedure described in [22]. The idea is to link the process (30) to a continuous time one by defining random durations for each of the iterations in (30), distributed according to

$$\text{Prob}[\ell \text{ steps at time } t \in \mathbb{R}^+] = \frac{1}{\ell!} \left[\frac{t}{\Delta_N} \right]^\ell e^{-t/\Delta_N}$$

The parameter Δ_N gives the average real-time duration of a single iteration of (30). Upon writing the Markov process (30) in probabilistic form[¶],

$$p_{\ell+1}(\mathbf{q}) = \int d\mathbf{q}' W(\mathbf{q}|\mathbf{q}') p_\ell(\mathbf{q}') \quad (31)$$

we find the new continuous-time process to be defined as

$$\frac{d}{dt} p_t(\mathbf{q}) = \frac{1}{\Delta_N} \left\{ \int d\mathbf{q}' W(\mathbf{q}|\mathbf{q}') p_t(\mathbf{q}') - p_t(\mathbf{q}) \right\} \quad (32)$$

The price paid is uncertainty about where one is on the time axis. This uncertainty, however, will vanish if we choose $\Delta_N \rightarrow 0$ for $N \rightarrow \infty$. Upon studying the scaling properties with N of the process (32) via the Kramers-Moyal expansion, for the kernel $W(\mathbf{q}|\mathbf{q}')$ corresponding to (30), one can establish the following facts:

- The correct scaling for the parameters Δ_N and η is: $\Delta_N = \mathcal{O}(N^{-1})$, $\eta = \mathcal{O}(N^0)$
- As $N \rightarrow \infty$ one obtains a Fokker-Planck equation for individual components q_i :

$$\frac{d}{dt} p_t(q_i) = \frac{\partial}{\partial q_i} [p_t(q_i) F_i(q_i, t)] + \frac{1}{2} \frac{\partial^2}{\partial q_i^2} [p_t(q_i) D_i(q_i, t)] + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$$

- The process for $\mathbf{q} = (q_1, \dots, q_N)$, however, cannot for $N \rightarrow \infty$ be described by a Fokker-Planck equation (non-Gaussian fluctuations in the q_i cannot a priori be discarded; although small, they might conspire for an N -component system).

Working out the diffusion matrix $D_{ij}(\mathbf{q})$ of the full process, for additive decision noise (23), for comparison with previous results in [17] and [18] (where the MG was claimed to obey a Fokker-Planck equation), reveals that $D_{ij}(\mathbf{q}) = D_{ij}^A(\mathbf{q}) + D_{ij}^B(\mathbf{q})$, where

$$D_{ij}^A(\mathbf{q}) = \frac{1}{N} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu \left[\Omega_\mu + \frac{1}{\sqrt{N}} \sum_j \xi_j^\mu \tanh(\beta q_j) \right]^2 \quad (33)$$

[¶] We leave aside the external fields for now, which only serve to define response functions.

$$D_{ij}^B(\mathbf{q}) = \frac{1}{N^2} \sum_{\mu=1}^p \xi_i^\mu \xi_j^\mu \sum_k (\xi_k^\mu)^2 [1 - \tanh^2(\beta q_k)] \quad (34)$$

The term $D_{ij}^A(\mathbf{q})$ mainly describes fluctuations due to the randomly drawn indices $\mu(\ell)$; $D_{ij}^B(\mathbf{q})$ reflects mainly the fluctuations generated by decision noise. Comparison shows that the diffusion matrices given in [17] and [18] are both approximations.

Having resolved the temporal regularization, one can now proceed to analyze the MG upon adapting to continuous time each of the generating functional analysis steps taken for the batch MG models (upon adding external fields). Since we choose $\Delta_N = \mathcal{O}(N^{-1})$ the impact of studying (32) rather than (31) (i.e. the uncertainty of our clock) is guaranteed to vanish as $N \rightarrow \infty$. The disorder-averaged generating functional is now a path integral:

$$\overline{Z[\boldsymbol{\psi}]} = \overline{\langle e^{-i \sum_i \int dt \psi_i(t) q_i(t)} \rangle} \quad (35)$$

Accepting the potential technical problems associated with steepest descent integration for path-integrals (which here imply additional assumptions on the smoothness of the correlation- and response functions), one can adapt the batch procedures, take the limit $N \rightarrow \infty$, and find again an effective ‘single agent’ process:

$$\frac{d}{dt} q(t) = \theta(t) - \alpha \int_0^t dt' R(t, t') \langle \sigma[q(t'), z] \rangle_z + \sqrt{\alpha} \eta(t) \quad (36)$$

The retarded self-interaction kernel is given by

$$R(t, t') = \delta(t - t') + \sum_{\ell > 0} (-1)^\ell (G^\ell)(t, t') \quad (37)$$

(i.e. the continuous-time equivalent of the previous $[\mathbf{I} + G]^{-1}$, written as a power series), whereas the zero-average Gaussian noise $\eta(t)$ is here characterized by

$$\begin{aligned} \langle \eta(t) \eta(t') \rangle &= \sum_{\ell, \ell' \geq 0} (-1)^{\ell + \ell'} \int_0^\infty ds_1 \dots ds_\ell ds'_1 \dots ds'_{\ell'} \prod_{uv} \left[1 + \frac{1}{2} \eta \delta(s_u - s'_v) \right] \\ &\times G(s_0, s_2) \dots G(s_{\ell-1}, s_\ell) [1 + C(s_\ell, s'_\ell) G^\dagger(s'_\ell, s'_{\ell-1}) \dots G^\dagger(s'_1, s'_0)] \end{aligned} \quad (38)$$

(which would only for $\eta = 0$ be a continuous-time version of previous expressions). The order parameter kernels C and G are to be solved from

$$C(t, t') = \langle \langle \sigma[q(t), z] \rangle_z \langle \sigma[q(t'), z] \rangle_z \rangle \quad (39)$$

$$G(t, t') = \frac{\delta}{\delta \theta(t')} \langle \langle \sigma[q(t), z] \rangle_z \rangle \quad (40)$$

The closed macroscopic equations (36,37,38,39,40) for the original so-called on-line MG with random external information (in the limit $N \rightarrow \infty$) are claimed to be exact for any α and any type of decision noise. Although they are again not easily solved in general, they have clarified many of the debates and apparent contradictions in previous studies. More specifically, they reveal that

- The ergodic stationary state in the regime $\alpha > \alpha_c(T)$, which is calculated easily, is *independent* of η , and consequently identical to that of the batch MG and the replica solution in [11, 12]. For $\alpha < \alpha_c(T)$, however, this is no longer the case: batch and on-line MG dynamics are now truly different. Previous methods had not yet succeeded in producing exact macroscopic laws for this regime.
- Premature truncation of the Kramers-Moyal expansion after the diffusion term (leading to a Fokker-Planck equation), provided with the correct diffusion matrix (i.e. neither of the ones proposed so far), would for $N \rightarrow \infty$ have led to the correct macroscopic equations. Apparently, the natural order parameters of the problem are not sensitive to the weak non-Gaussian fluctuations.

7. Discussion

Minority Games are a specific class of disordered stochastic systems without detailed balance. In this contribution I have tried to give a (biased) introduction to their definitions, properties, and non-equilibrium statistical mechanical theories. I claim that the generating functional analysis techniques proposed by De Dominicis [4] (which allow for the derivation of exact, closed macroscopic dynamic equations in the limit $N \rightarrow \infty$) are the canonical tools with which to study these systems, both (a priori) in view of the specific nature of the MGs as statistical mechanical systems and (retrospectively) in view of the insight and clarification which application of these techniques have already generated. Compared to some of the other methods in the disordered systems arena, generating functional analysis is immediately seen to be in principle mathematically sound. As always in applied areas, one should worry about the existence of limits and integrals, but these are quite conventional reservations (in contrast to e.g. replica theory); the only serious step is the application of saddle-point arguments in path integrals (for the on-line MG), but, again, this is a familiar type of worry in statistical mechanics about which much has already been written.

In addition the generating functional analysis methods have opened the door to many interesting new questions which can now start to be addressed. At the level of increasing realism, one could re-introduce multiple strategies ($S > 2$), increase the agents' decision range (e.g. to $\sigma \in \{-1, 0, 1\}$, representing the options buy, inactivity, or sell), inspect slowly evolving strategy tables, consider the trading of multiple commodities, introduce capital and trading costs, etc. This is perhaps not likely to generate fundamentally new mathematical puzzles, but would make MGs more relevant to those interested in markets as such, and thereby promote the power of non-equilibrium statistical mechanics in a wider community.

At the mathematical level one could now assess the precise impact of replacing the true market memory by random numbers in the definition of the external information (there is simulation evidence that for $\alpha < \alpha_c$ there might be non-negligible differences after all). In replica approaches this would be nearly impossible (since it renders the model again non-Markovian), but within the generating functional framework (where

one already finds a non-Markovian effective ‘single agent’ anyway) there appear to be no fundamental obstacles. Secondly, one could try (for the batch and on-line versions) to construct or approximate further solutions of the macroscopic equations in the non-ergodic regime $\alpha < \alpha_c$. Thirdly, the generating functional equations reveal that many of the approximations which have been used in the past to calculate the volatility σ boil down to assuming the existence of non-equilibrium Fluctuation-Dissipation Theorems in the stationary state. It would be interesting to investigate more generally the existence and nature of such identities relating the correlation and response functions for models such as the MG, where (in contrast to e.g. spin-glasses) detailed balance violations are built-in microscopically, rather than emerging only in the $N \rightarrow \infty$ limit. Any new results concerning such identities would have considerable implications for a wider area of non-equilibrium systems, including e.g. non-symmetric neural networks, cellular automata, and immune system networks.

Acknowledgements

It is my pleasure to thank my collaborators Alexander Heibel and David Sherrington.

References

- [1] Challet D and Zhang Y-C 1997 *Physica A* **246** 407-418
‘Emergence of cooperation and organization in an evolutionary game’
- [2] Arthur W B 1994 *Am. Econ. Assoc. Papers and Proc.* **84** 406-411
‘Inductive reasoning and bounded rationality’
- [3] Challet D, <http://www.unifr.ch/econophysics/minority>
‘Minority Game’s web page’
- [4] De Dominicis C 1978 *Phys. Rev. B* **18** 4913-4919
‘Dynamics as a substitute for replicas in systems with quenched random impurities’
- [5] Cavagna A 1999 *Phys. Rev. E* **59** R3783-R3786
‘Irrelevance of memory in the minority game’
- [6] Savit R, Manuca R and Riolo R 1999 *Phys. Rev. Lett.* **82** 2203-2206
‘Adaptive competition, market efficiency and phase transitions’
- [7] Cavagna A, Garrahan J P, Giardina I, and Sherrington D 1999 *Phys. Rev. Lett.* **83** 4429-4432
‘Thermal model for adaptive competition in a market’
- [8] Hart M, Jefferies P and Johnson N F 2000 *Phys. Rev. E* **63** 017102
‘Generalized strategies in the minority game’
- [9] Jefferies P, Hart M, Johnson N F and Hui P M 2000 *J. Phys A: Math. Gen.* **33** L409-L414
‘Mixed population minority game with generalized strategies’
- [10] Challet D and Marsili M 1999 *Phys. Rev. E* **60** R6271-R6274
‘Symmetry breaking and phase transition in the minority game’
- [11] Challet D, Marsili M and Zecchina R 2000 *Phys. Rev. Lett.* **84** 1824-1827
‘Statistical mechanics of heterogeneous agents: minority games’
- [12] Marsili M, Challet D and Zecchina R 2000 *Physica A* **280** 522-553
‘Exact solution of a modified El Farol’s bar problem: efficiency and the role of market impact’
- [13] Mézard M, Parisi G and Virasoro MA 1987 *Spin-Glass Theory and Beyond* (Singapore: World Scientific)

- [14] Kühn R, Bös S and van Hemmen JL 1991 *Phys. Rev. A* **43** 2084-2087
‘Statistical mechanics for networks of graded-response neurons’
- [15] Challet D, Marsili M and Zecchina R 2000 *Phys. Rev. Lett.* **85** 5008-5008
‘Comment on Thermal Model for Adaptive Competition in a Market’
- [16] Cavagna A, Garrahan J P, Giardina I and Sherrington D 2000 *Phys. Rev. Lett.* **85** 5009-5009
‘Cavagna et al reply’
- [17] Garrahan J P, Moro E and Sherrington D 2000 *Phys. Rev E* **62** R9-R12
‘Continuous time dynamics of the minority game’
- [18] Marsili M and Challet D 2001 *Phys. Rev. E* **64** 056138
‘Continuum time limit and stationary states of the minority game’
- [19] Heibel J A F and Coolen A C C 2001 *Phys. Rev. E* **63** 056121
‘Generating functional analysis of the dynamics of the batch minority game with random external information’
- [20] Coolen A C C, Heibel J A F and Sherrington D 2001 *Phys. Rev. E* **65** 016126
‘Dynamics of the batch minority game with inhomogeneous decision noise’
- [21] Coolen A C C and Heibel J A F 2001 *J. Phys A: Math. Gen.* **34** 10783-10804
‘Dynamical solution of the on-line minority game’
- [22] Bedeaux D, Lakatos-Lindenberg K and Shuler K 1971 *J. Math. Phys.* **12** 2116-2123
‘On the relation between master equations and random walks and their solutions’