On non-ergodic phases in Minority Games

Andrea De Martino[(1)] and Tobias Galla[(2)]

[(1)]INFM/CNR and Dipartimento di Fisica
Università di Roma "La Sapienza", Rome, Italy
[(2)]The Abdus Salam International Centre for Theoretical Physics,
Strada Costiera 11, 34014 Trieste, Italy

Abstract

In the first part of this review, we survey the known analytic results regarding the non-ergodic phase of standard Minority Games. Such phases are characterized by the fact that, at odds with ergodic regimes, the steady state properties of the game (e.g. the volatility) depend both on the initial conditions chosen for the agents' learning process as well as on the learning rate. Secondly, we present a discussion of the effects of finite-memory learning, which lifts the non-ergodicity, in the context of spherical Minority Games.

1 Introduction

The Minority Game (MG, [1]) is perhaps the simplest agent-based model that is able, in certain variants, to reproduce some of the empirically observed regularities that characterize the behavior of price returns in financial markets (for a recent review focused on the MG as a market model see [2]). It is a toy model, in the sense that the dynamical laws governing the choices of agents are considerably stylized and thus analytically tractable to a high degree, provided one recognizes some important similarities with models of neural networks and disordered magnetic alloys (both well-known in the physics of disordered systems). For this reason, it has drawn much interest from statistical physicists. Several monographs and reviews (as well as web resources) deal with the details of the model and we refer the reader to them for a more thorough introduction and motivation [3–7].

At the core of the model lies the assumption that at each time step N agents react to the receipt of one of P possible information patterns by either buying or selling, in the attempt to anticipate the minority (buy low/sell high). Agents who take the correct minority decision (i.e. those who buy when the majority is selling and vice versa) are rewarded. Interaction between the agents is only indirect, that is players cannot identify the precise actions of other agents, but react only to the aggregate action of all players in the market (the price and its movements). Each agent holds a pool of strategies based on which he makes his trading decisions mapping the publicly available piece of information onto a binary decision (i.e. whether to buy or to sell), and aims at identifying his best strategy using virtual scores to monitor the performance of a particular strategy in the past. At each round of the game every agent updates the score of each of his strategies: the scores of strategies which would have predicted the correct minority decision are increased while those of strategies which yield unprofitable trading decisions are reduced.

The emerging picture, which persists through almost all studied variants, is that, in the limit where $N \to \infty$, the overall behavior changes drastically when the (finite) parameter $\alpha = P/N$ crosses a critical value α_c. Specifically, the volatility of the total bid undergoes a non-trivial phase transition. The bid here serves as a proxy for the price return in real market time series, so that the volatility reflects the magnitude of price fluctuations. In the supercritical phase ($\alpha > \alpha_c$), the dynamics is ergodic, that is the stationary state, which can be characterized through a few macroscopic quantities like the volatility, is unique and independent of the initial conditions of the agents' learning process (their 'prior beliefs'). In the physical literature, this phase is often called 'asymmetric', in reference to the fact that an explicit symmetry of the model (equal a priori probability of buying or selling) appears to be broken in the steady state, where one action is performed more frequently than the other in the presence of some information patterns. Note that this implies that the price signal possesses some predictability. In the subcritical phase ($\alpha < \alpha_c$), instead, one encounters many steady states depending on the initial conditions. Stated differently, the dynamics breaks ergodicity below the transition. Remarkably, both very efficient (low volatility) and very inefficient (high volatility) states can emerge with a slight change of the agents' prior beliefs. Contrary to the supercritical phase, the non-ergodic regime is symmetric, that is the symmetry between buying and selling is preserved asymptotically[1].

While the behavior of the model in the asymmetric phase has been studied in great mathematical detail, much less is known about the symmetric phase. Here the breakdown of ergodicity is fatal for all known analytical techniques and one is forced to look for *ad hoc* solutions to every specific issue that may arise. Thus, only a few problems have been addressed so far. Understanding the symmetric phase is important for several reasons, including some practical ones: indeed MGs suggest that real markets operate close to the phase transition, on the edge between ergodicity and non-ergodicity.

We start this note by defining the model (Sec. 2) and pass then to reviewing the known properties of the symmetric phase of the standard MG (Secs. 3 and 4), considering also the role played by market impact (Sec. 5). Next we shall address the role of finite memory in the learning dynamics in a context where it was previously not studied, that of the so-called 'spherical' MGs (Sec. 6). Our focus will be set mainly on the physical ideas behind the different approaches rather than on the technical details, in the hope to emphasize the overall structure of the model's dynamics. Indeed, while strong non-ergodicity may be enough to prevent a full analytic treatment, a phenomenological theory would be highly desirable and we believe that a physical understanding is necessary in order to develop such a theory.

2 The model

The MG describes a system of N agents which will be labelled with Roman indices i, j, and proceeds in discrete rounds (trading periods) $t = 0, 1, 2....$ At each time step t every agent i takes a trading decision $b_i(t) \in \mathbb{R}$ (a 'bid') in response to the observation of a public information pattern $\mu(t)$. In the original version of the MG this information represented the actual market history [1, 10–15], but it has been

[1] It is worth remarking that some variants of the MG do not exhibit non-ergodic symmetric phases. Examples can be found in [8, 9].

shown that results do not change qualitatively if the information patterns are drawn at random at each step [16]. What is important, however, is that all agents are presented with the same information at each time step, as this induces the effective interaction and co-ordination between agents. We will in the following consider the MG with random external information, and will assume that $\mu(t)$ is chosen randomly and independently from a set with $P = \alpha N$ possible patterns, i.e. $\mu(t) \in \{1, \ldots, \alpha N\}$. One then defines the so-called re-scaled total market bid at round t as $A(t) = N^{-1/2} \sum_i b_i(t)$ (the re-scaling factor $N^{-1/2}$ here simplifies the subsequent theoretical analysis, as it renders $A(t)$ finite in the so-called thermodynamic limit $N \to \infty$).

To take his trading actions, each agent i holds a pool of S fixed trading strategies (look-up tables) $\mathbf{R}_{ia} = (R_{ia}^1, \ldots, R_{ia}^P)$, with $a = 1, \ldots, S$. Typically we will here focus on the case $S = 2$, as the behaviour of the model does not change qualitatively for a larger number of strategies per player [17, 18]. If agent i decides to use strategy a in round t of the game, his trading action at this stage will be $b_i(t) = R_{ia}^{\mu(t)}$. All strategies $\mathbf{R}_{ia}$ are chosen randomly and without correlation before the dynamics is started; they represent the agents' heterogeneity, in that they are generically expected to react differently from each other (in physics jargon, this is the "quenched disorder" of the problem). The standard choice in the literature are binary strategy entries, $\mathbf{R}_{ia} \in \{-1, 1\}^P$, but modification to continuous entries or MGs with inner product definitions have been considered [10, 15].

In order to decide which strategy to use the agents keep track of a score $p_{ia}(t)$ which they allocate to each of their strategies, so that $p_{ia}(t)$ denotes the score of agent i's strategy table number a at time t. These score valuations are based on the success had the player always played that particular strategy, i.e one has

$$p_{ia}(t+1) = p_{ia}(t) - R_{ia}^{\mu(t)} A(t). \tag{1}$$

The minus sign ensures that strategies which would have produced a minority decision are rewarded (in this case $R_{ia}^{\mu(t)} A(t)$ is negative, so that the score p_{ia} is increased). At each round t a given player i then uses the strategy in his pool with the highest score, i.e. $b_i(t) = R_{i\tilde{a}_i(t)}^{\mu(t)}$, where $\tilde{a}_i(t) = \arg\max_a p_{ia}(t)$. For $S = 2$ the rules (1) can be simplified as only the difference between the two scores,

$q_i(t) = \frac{1}{2}[p_{i1}(t) - p_{i2}(t)]$ is relevant for player i's actions: he plays strategy $\mathbf{R}_{i1}$ in round t if $q_i(t) > 0$, and $\mathbf{R}_{i2}$ if $q_i(t) < 0$ (for this reason, $q_i(t)$ is sometimes referred to as the 'preference' of agent i at time t). Player i's bid in round t then reads $b_i(t) = \omega_i^{\mu(t)} + \mathrm{sgn}[q_i(t)]\xi_i^{\mu(t)}$, where $\boldsymbol{\omega}_i = \frac{1}{2}[\mathbf{R}_{i1} + \mathbf{R}_{i2}]$ and $\boldsymbol{\xi}_i = \frac{1}{2}[\mathbf{R}_{i1} - \mathbf{R}_{i2}]$. The above update rule for the scores $\{p_{ia}(t)\}$ can then be cast as the following update prescription for the score differences $\{q_i(t)\}$:

$$q_i(t+1) = q_i(t) - \xi_i^{\mu(t)} \left[\Omega^{\mu(t)} + \frac{1}{\sqrt{N}} \sum_j \xi_j^{\mu(t)} \mathrm{sgn}[q_j(t)] \right], \tag{2}$$

with $\boldsymbol{\Omega} = N^{-1/2} \sum_j \boldsymbol{\omega}_j$. Equation (2) defines the standard so-called 'on-line' MG. Numerical simulations of the MG have established that only one basic parameter controls the dynamics of the model, namely the ratio $\alpha = P/N$ of the number of different information patterns over the number of players in the system. I.e. differently sized populations of agents behave quantitatively identically if the number of values P the information patterns can take is re-scaled as to keep P/N constant. The key observables of the MG are the so-called volatility σ^2 and the predictability H of the resulting time-series of bids. While the predictability is given by

$$H = \frac{1}{P} \sum_{\mu=1}^{P} \langle A|\mu \rangle_t \tag{3}$$

with $\langle \cdot|\mu \rangle_t$ an time-average conditioned on the occurence of information pattern μ, σ^2 is simply defined as the variance of the time series $A(t)$

$$\sigma^2 = \left\langle A^2(t) \right\rangle_t \tag{4}$$

$H = 0$ indicates that $\langle A|\mu \rangle_t = 0$ for all μ, i.e. the time-series $A(t)$ is unpredictable. If $H > 0$, however, there exist μ so that the sign and magnitude of the bid $A(t)$ given a state μ can be predicted in a statistical sense. The volatility σ^2 is instead a measure for the global performance of the market. If $\sigma^2 \approx 0$ then buying and selling bids are roughly matched at each time step (a perfect balance corresponds to $A(t) = 0$). In this case of (roughly) equally many buyers and sellers a large fraction of agents will be in the minority at any given time-step (imagine a situation in which

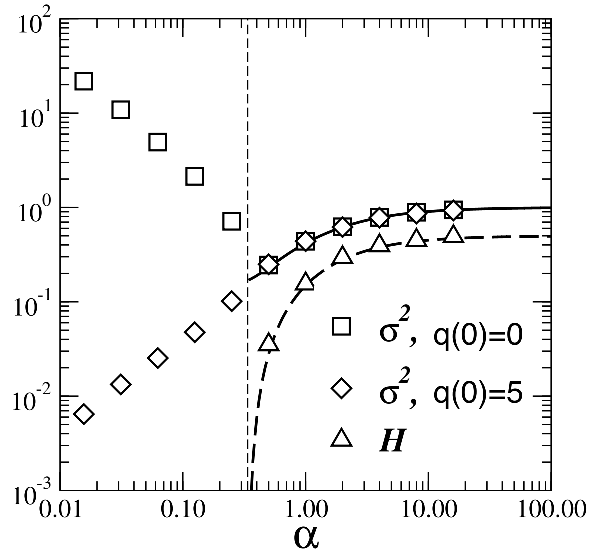

Figure 1: σ^2 and H as a function of α for different initial conditions ($q_i(0) = q(0)$ for all i at the start of the game). The two branches of σ^2 for small α behave respectively as $1/\alpha$ (top) and α (bottom). See Eq. (15).

50 agents play $+1$ and 51 agents play -1 at a given step, resulting in 50 winners and 51 losers), hence the number of 'winners' among the N agents takes its maximal value $N/2$. The population of agents is globally successful. If however σ^2 is large, then supply and demand are not well matched, and a only a small fraction of the agents win at any step (one may for example think of a situation in which e.g. 91 out of $N = 101$ agents play $+1$, while the other 10 play -1. Then only 10 out of 101 agents are winners in this particlar round). A more careful analysis shows indeed that the overall payoff cumulated over the entire population is given by $-\sigma^2$ (the overall payoff is always non-positive, as the MG is intrinsically a negative-sum game). Note that the profit of each winner is smaller the smaller is σ^2, because it has to be shared with a larger number of other winners.

Plotting σ^2 and H as a function of α (see Fig. 1) reveals several interesting features of the standard MG:

- A phase with vanishing predictability $H = 0$ is observed for small $\alpha = P/N$, separated from one with positive predictability at large α. The two regimes are separated by a phase transition at $\alpha_c = 0.3374...$

- The transition point between these two phases corresponds to a minimum in the volatility σ^2 if the dynamics is started from so-called *tabula rasa* initial conditions ($q_i(0) = 0$ for all i).

- Above α_c initial conditions have no influence on the stationary states of the dynamics. Neither H nor σ^2 depend on the starting point. Below α_c however different branches for σ^2 are observed for different biases $q_i(0)$ of the initial score-valuations. In particular one finds solutions in which $\sigma^2 \to 0$ for $\alpha \to 0$, and others in which the volatility diverges as the limit $\alpha \to 0$ is approached.

The latter phase at low α is referred to as the symmetric phase ($H = 0$), or as the *non-ergodic* phase in the literature, as the stationary state is here strongly dependent on initial conditions. The remainder of this review will focus mostly on the behaviour of the MG in this non-ergodic phase at small $\alpha = P/N < \alpha_c$. We here note that similar transitions are observed ina a variety of extensions and modifications of the MG, but that the numerical value of α_c can depend on the specific model considered.

3 The dynamical approach: role of initial conditions

To simplify the dynamical analysis of the MG process, it is common to address the so-called 'batch limit', first introduced in [19], where one considers the dynamics in terms of an average over all possible values of the external information in (2). This corresponds to updating the $\{q_i\}$ only once every $\mathcal{O}(N)$ time-steps, and leads to the the so-called 'batch' dynamics[2]:

$$q_i(t+1) = q_i(t) - \frac{2}{N} \sum_j \left[\sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu s_j(t) + \sum_{\mu=1}^{P} \xi_i^\mu \omega_j^\mu \right]. \tag{5}$$

[2]A detailed analysis shows that the phase behaviour of batch and on-line games is identical, and only small quantitative differences are observed for the volatility. See however [20] for versions of MGs where on-line and batch MGs differ in their qualitative behaviours.

Here $s_j(t) = \text{sgn}[q_j(t)]$. The dynamical approach employs path-integrals to calculate relevant macroscopic observables directly from the dynamics. The net result of the (lengthy) procedure is that the set of N globally coupled processes above can be re-cast in the form of a single stochastic process possessing the same steady state properties. It is customary to think that such process describes the behaviour of a single (representative) trader. This process is given by

$$q(t+1) = q(t) - \alpha \sum_{t' \leq t} (\mathbb{1} + G)^{-1}_{tt'} s(t') + \sqrt{\alpha}\eta(t), \tag{6}$$

where $s(t) = \text{sgn}[q(t)]$ and $\eta(t)$ is a coloured Gaussian noise with covariance matrix as follows

$$\left\langle \eta(t)\eta(t') \right\rangle = [(\mathbb{1} + G)^{-1} D (\mathbb{1} + G^T)^{-1}]_{tt'}. \tag{7}$$

The functions $C_{tt'}$ and $G_{tt'}$, called respectively correlation and response functions, are to be computed self-consistently from

$$C_{tt'} = \left\langle s(t)s(t') \right\rangle, \qquad G_{tt'} = \frac{1}{\sqrt{\alpha}} \frac{\partial}{\partial \eta(t')} \left\langle s(t) \right\rangle, \tag{8}$$

where $\langle \cdot \rangle$ denotes an average over realisations of the effective process, and where $D_{tt'} = 1 + C_{tt'}$ for all t, t'. While the physical meaning of the correlation function $C_{tt'}$ is clear, that of $G_{tt'}$ is more subtle: it measures how much the representative agent's decision would change at time t due to a small perturbation (or a change of the noise) applied at time t'. Practically, it quantifies how sensitive agents are to small perturbations. The main aspect of this procedure to be kept in mind is that, at odds with the original set of equations, the representative process is noisy. Moreover, the noise is correlated in time, and the process is non-Markovian as the second term on the right-hand-side of (6) represents a sum over all past time steps $t' \leq t$.

3.1 The ergodic regime

In the ergodic regime ($\alpha > \alpha_c$) initial conditions have no influence on the dynamics in the long-run. An analytical solution is here possible based on a so-called time-translation invariance ansatz $C_{tt'} = C(t-t')$, $G_{tt'} = G(t-t')$. Furthermore ergodicity is reflected by a finite integrated response $\chi = \int dt G(t) < \infty$, as perturbations decay fast enough. With these self-consistent assumptions explicit equations for the integrated resonse χ and the persistent part c of the correlation function can be found. Specifically one has [19]

$$c = 1 - \left(1 - \frac{1+c}{\alpha}\right) \operatorname{erf}\left(\sqrt{\frac{\alpha}{2(1+c)}}\right) - 2\sqrt{\frac{1+c}{2\pi\alpha}} e^{-\alpha/(2(1+c))} \tag{9}$$

and

$$\chi^{-1} = \frac{\alpha}{\operatorname{erf}\left(\sqrt{\frac{\alpha}{2(1+c)}}\right)} - 1. \tag{10}$$

In the derivation of these equations it turns out to be crucial that agents can be divided into two classes: (i) so-called frozen agents, who in the long-run never change strategies, (ii) so-called fickle agents. These are all other non-frozen agents, i.e. those who always or occasionally change between their two look-up tables.

We will not enter the further details here, but will only say that the predicatability H can be obtained exactly from the solutions of the two equations given above as

$$H = \frac{1}{2} \frac{1+c}{(1+\chi)^2} \tag{11}$$

and that the volatility can be approximated as[3]

$$\sigma^2 = \frac{1}{2} \frac{1+c}{(1+\chi)^2} + \frac{1}{2}(1-c) \tag{12}$$

These results are valid in the ergodic phase only, as they rely on the time-translation invariant ansatz and on a sufficiently quickly decaying response function with finite integral ($\chi < \infty$). The ergodic theory hence predicts its own breakdown self-consistently as the point at which χ diverges. Numerical solution of the above equations reveals that this occurs when $\alpha \downarrow \alpha_c$, where $\alpha_c = 0.3374....$

3.2 The non-ergodic regime: limit $\alpha \to 0$

Studying the dynamics in the non-ergodic phase is much more intricate, as initial conditions here matter crucially. This complicates the theoretical analysis as one

[3]There are several different approximate expressions in the literature (see e.g. [4] for details). We here only give the most common one.

needs to keep track of the starting point of dynamics, and may not restrict the attention to a stationary state. As we have seen in Fig. 1 several branches of solutions exists below the phase transition, which one of these is reached by the dynamics asymptotically depends on how the system is initialised, e.g. on the initial bias in the strategy valuations of players.

Analytical progress is however possible in the limit of only few information patterns P, more precisely in the limit $P \ll N$, i.e. $\alpha \to 0$. In this regime one observes in simulations that all agents belong to one of the two following groups:

(i) Frozen agents: these are, as above, agents who asympotically never switch strategies, i.e they have $s_i(t) \equiv +1$ or $s_i(t) \equiv -1$ independently of t.

(ii) Oscillatory agents: these are agents who switch strategies at every time step, i.e for such agents one has $s_i(t) = (-1)^t$ or all t or $s_i(t) = -(-1)^t$ for all t.

This observation allows one to make an ansatz of the form $C_{t+\tau,t} = c + (1-c)(-1)^\tau$ for the correlation function, and $\eta(t) = \gamma z\sqrt{(1-c)}(-1)^t$ for the single-agent noise $\eta(t)$, with $\gamma \equiv [1+\widetilde{G}(\pi)]^{-1} = \sum_t (-1)^t (\mathbb{1}+G)^{-1}(t)$, where $\widetilde{G}(\omega)$ denotes the Fourier transform of the time-translation invariant response function $G(t)$. z in the above expression is a static Gaussian variable of zero mean and unit variance, and reflects the stochasticity of the representative agent process within the simplications of the ansatz used. Is is then possible to find explicit ansätze also for the trajectories of frozen and oscillatory representative agents. We will not give details here, but refer to [4, 19] for their precise algebraic forms. Upon using these ansätze one then finds the following equations for the volatility σ^2 and the fraction of frozen agents ϕ in the limit $\alpha \to 0$ [19]:

$$\sigma = \int dq P(q) \frac{\exp\left(-q^2/(\alpha\sigma^2)\right)}{\sqrt{\alpha\pi}} \tag{13}$$

$$\phi = \int dq P(q) \mathrm{erf}\left(\frac{|q|}{\sigma\sqrt{\alpha}}\right) \tag{14}$$

The problem is hence reduced to finding the distribution $P(q)$ of score-differences. This quantity contains all remaining memory of initial conditions. Assuming it takes a Gaussian form $P(q) \sim e^{-(q-q_0)^2/(2\Lambda^2)}$ with $q_0 \equiv q(0)$ the bias imposed initially further information about the volatility σ can be extracted in the limit $\alpha \to 0$. Namely one finds high volatility solutions for sufficiently small initial biases, and low-volatility solutions for large initial bias:

$$\alpha \to 0 \quad : \qquad \sigma^2 \sim \begin{cases} 1/\alpha & \text{for } q_0 < q_{0c} \equiv (2\pi e)^{-1} \\ \alpha & \text{for } q_0 > q_{0c} \end{cases} \tag{15}$$

Note that the critical value q_{0c} can here be obtained analytically as $q_{0c} = 1/(2\pi e)$ [19].

4 The static approach: the role of the learning rate

We have seen that the dynamics of preferences in the case $S = 2$ takes the simple form

$$q_i(t+1) - q_i(t) = -\overline{\xi_i A(t)} \quad , \quad A(t) = \Omega^\mu + \frac{1}{\sqrt{N}} \sum_i s_i(t) \xi_i^\mu, \tag{16}$$

where the over-line denotes the 'batch' averaging over μ's. Let us generalize the choice rule $s_i(t) = \text{sgn}[q_i(t)]$ to the logit form

$$\text{Prob}\{s_i(t) = s\} \propto e^{\Gamma p_{is}(t)} \tag{17}$$

(we have here labeled the two strategies of player i by $s = \pm 1$ to compactify the notation). A finite value of Γ introduces decision noise (the original deterministic update rule, in which the better-performing strategy is chosen stricly at each step, corresponds to $\Gamma \to \infty$). At finite Γ it is possible to derive the continuous-time equivalent of Eq. (16) for the re-scaled variable $y_i(t) = \Gamma q_i(t)$. The procedure is detailed in [6, 21]. The resulting Langevin equation contains a noise term which accounts for the volatile decision dynamics:

$$\frac{dy_i}{dt} = -\overline{\xi_i \Omega} - \sum_j \overline{\xi_i \xi_j} \tanh y_i + z_i(t), \tag{18}$$

$$\langle z_i(t) \rangle = 0, \quad \langle z_i(t) z_j(t') \rangle = \frac{\Gamma \sigma^2}{P} \overline{\xi_i \xi_j} \delta(t - t'). \tag{19}$$

This means that in the continuous-time limit the preferences are subject to stochastic fluctuations of magnitude $\sqrt{\Gamma}$ around their average, which evolves according to

$$\frac{d\langle y_i\rangle}{dt} = -\overline{\xi_i\Omega} - \sum_j \overline{\xi_i\xi_j}\langle\tanh y_i\rangle \tag{20}$$

One sees that, perhaps counter-intuitively, for $\Gamma \to 0$ the dynamics becomes deterministic and admits the Lyapunov function

$$H = \frac{1}{P}\sum_\mu \left[\Omega^\mu + \sum_i \xi_i^\mu m_i\right]^2 \tag{21}$$

with $m_i = \langle\tanh y_i\rangle$. In other words, the stationary states are described by the minima of H in $[-1,1]^N$ (the domain of variability of the m_i's). This makes the $\Gamma \to 0$ case particularly simple to analyze. Indeed, the minimization of H can be carried out using spin-glass techniques (see e.g. [3, 6, 13] for detailed accounts). Generally speaking, H is a non-negative definite quadratic form, implying that it attains its minimum on a connected set of point (eventually a single point). Indeed in the ergodic phase ($\alpha > \alpha_c$), it turns out that H has a single point-like minimum, where it has a non-zero (positive) value (we shall denote the minimum value of H by $H_\star$). As a consequence, the steady state is unique. Now $H_\star$ decreases as $\alpha \downarrow \alpha_c$ and vanishes at and below α_c (non-ergodic phase). Below α_c, the equation $H = 0$ admits many solutions (it roughly corresponds to P conditions of the form $\Omega^\mu + N^{-1/2}\sum_i \xi_i^\mu m_i = 0$ which have to be satisfied by N variables with $N > P$), which as said above must form a connected subset of $[-1,1]^N$. This implies the existence of multiple steady states and naturally a dependence on initial conditions.

When $\Gamma > 0$ the situation changes because H is no longer an exact Lyapunov function. From the viewpoint of individual agents, a larger Γ implies larger fluctuations of the preferences around their averages and thus a longer time needed to wash them out. However, it can be shown that only the non-ergodic phase is affected by a non-zero Γ. Its effects are particularly strong on the volatility $\sigma^2 = \overline{\langle A^2\rangle}$, which after some algebraic manipulations can be written as

$$\sigma^2 = H + \sum_i \overline{\xi_i^2}(1 - m_i^2) + \sum_{i\neq j} \overline{\xi_i\xi_j}\langle(\tanh y_i - m_i)(\tanh y_j - m_j)\rangle \tag{22}$$

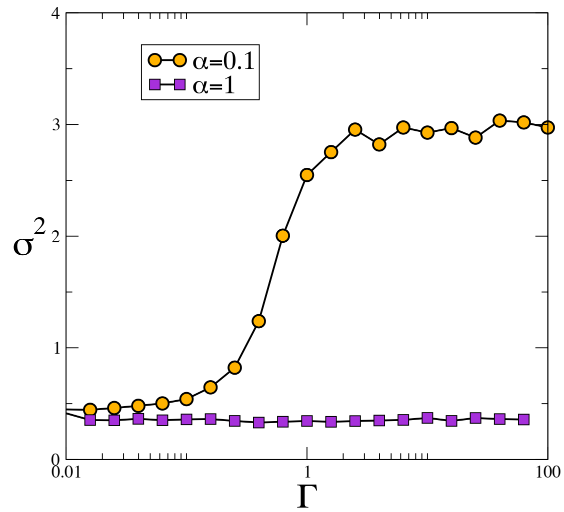

Figure 2: Volatility σ^2 versus Γ in the ergodic phase ($\alpha = 1$, lower curve) and in the non-ergodic phase ($\alpha = 0.1$, upper curve).

The last term on the right-hand-side is the problematic one. If agents i and j act in an uncorrelated way, then the average factorizes and, as $\langle \tanh y_i \rangle = m_i$ by definition, becomes identically zero. It can be shown [21] that below α_c this is not the case and one would have to evaluate the average in the last term explicitly. In particular, the average is performed with respect to the probability distribution of the preferences y_i's, which depend on Γ through the noise term in the Langevin equation. Hence if agents act in a correlated way a dependence on Γ is to be expected. Indeed numerically one finds that σ^2 increases with Γ when $\alpha < \alpha_c$, while the volatility is independent of Γ in the large-α phase (see Fig. 2). Below α_c, the dependence on Γ adds to that on initial conditions so that

$$\sigma^2 = \begin{cases} f(\alpha) & \text{for } \alpha > \alpha_c \\ g(\alpha, \Gamma, \{y_i(0)\}) & \text{for } \alpha < \alpha_c \end{cases} \tag{23}$$

where f and g are some functions. The former has been well approximated by different means. The latter is not known except for the case $\Gamma = 0$, $\{y_i(0) = 0\}$, when it can be recovered from the minimization of H. Note that in general the calculation of σ^2 requires solving a complicated self-consistent problem, since when it's not vanishing the average $\langle(\tanh y_i - m_i)(\tanh y_j - m_j)\rangle$ depends on σ^2 itself through the noise covariance in the Langevin equation (see (18)). However it has been shown that for $\{y_i(0) = 0\}$, below α_c, the solution for $\Gamma \ll 1$ can be expanded as [21]

$$\frac{\sigma^2}{N} = \frac{1-Q}{2}\left[1 + \frac{1-Q+\alpha(1-3Q)}{4\alpha}\Gamma + \mathcal{O}(\Gamma^2)\right] \tag{24}$$

where $Q = (1/N)\sum_i m_i^2$. This is to date the most accurate explicit result we have on the subcritical behavior of the volatility as a function of Γ.

5 The role of price impact

One of the key behavioral assumption of agent-based models of financial markets is that traders are price-takers, that is they neglect the impact they have on prices and treat them as an external signal (in a daily-life situation one might argue that a costumer who goes to a supermarket to buy -say- milk would not normally take into account that his buying milk might have an impact on the future price of milk, but that he treats the present and future prices as given quantities). The usual justification for this assumption is that it is normally impossible for a trader to estimate how much the price would have changed had he acted differently. Moreover, in a market with a large number N of traders all of which are similar in financial size, it is reasonable to assume that each trader's contribution will weigh $1/N$ in the aggregate leading to the price. So every individual impact is vanishingly small when $N \to \infty$. There are however some caveats. First, it is well known that a single trade can have a large impact on the price if for instance the market is in a state where few orders are present. This is clearly a crucial aspect to monitor if one is to develop automated trading systems. Secondly, it turns out that as soon as agents take even an infinitessimally small account of their impact, the market's global dynamical properties can be altered considerably.

In MGs, impact can be taken into account when each trader in his learning dynamics adjust the actual price return $A(t) = \sum_i b_i(t)$ by disentangling his bid and substituting it with the trading action he might have taken instead (e.g. buy instead of sell). If an agent were to take his impact on the price movements into account he would thus not use $A(t)$ in the update rule, but a modified global bid $A(t)-b_i(t)+R_{is}^{\mu(t)}$ (with $s \neq s_i(t)$), where his action ($b_i(t)$) is replaced by a potentially different one ($R_{is}^{\mu(t)}$), which he might have taken instead. This ensures that he estimates the profitability of strategy s using always the correct price return, namely that which would have occurred had he actually used s (see however [6] for a more thorough discussion). In order to allow for different degrees of agents' sophistication in the learning process, the following adjusted dynamics has been proposed [13]:

$$p_{ia}(t+1) - p_{ia}(t) = -R_{ia}^{\mu(t)} \left[A(t) - \eta \left(b_i(t) - R_{ia}^{\mu(t)} \right) \right] . \tag{25}$$

$\eta = 0$ and $\eta = 1$ correspond, respectively, to the original model (price-taking agents) and to the fully adjusted case ('sophisticated' traders). Varying η from 0 to 1 allows for an interpolation between these two limiting cases. This model has been studied both with dynamical and static methods [22, 23]. It is particularly interesting to analyze the phase structure of the model in the (α, η)-plane (see Fig. 3), the previous cases without impact correction corresponding to the line $\eta = 0$. One finds that the non-ergodic phase expands as η increases (i.e. the critical values of α for the onset of non-ergodicity become larger), until for $\eta = 1$ the ergodic phase disappears and the market is completely non-ergodic [4].

Perhaps surprisingly, for $\eta = 1$ the model can be simplified and more detailed information about the steady states can be drawn [22]. The simplification consists in the fact that all agents freeze, that is

$$\lim_{N\to\infty} \frac{1}{N} \sum_i \langle s_i \rangle_t^2 = 1 \qquad \text{for } \eta = 1 \tag{26}$$

at large times. In other words, asymptotically every trader locates an optimal strategy in his pool and sticks to it. This is the Nash equilibrium of the MG: no agent has an incentive to unilaterally switch strategies. As a consequence, agents here behave deterministically in the sense that upon receiving the same information pattern they

[4] We note however that the predictability H remains strictly positive for all α, when $\eta > 0$. The notion of non-ergodicity in MGs with impact correction is hence different from the one in the standard MG. The precise pattern of ergodicity breaking in such versions of the MG is discussed in [22, 23].

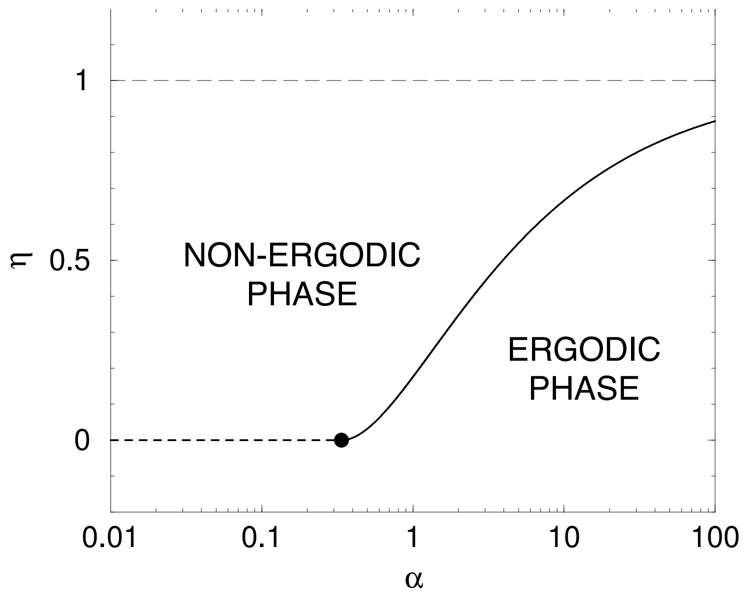

Figure 3: Phase structure of the MG with market impact correction.

always react in the same way. An intuitive reason for this is that for $\eta = 1$ agents no longer minimize a non-negative definite quadratic form (such as H). As a matter of fact, it can be shown that the quantity they minimize is precisely σ^2, which as a function of the m_i's is harmonic. This implies that it attains its minima on the corners of its configuration space, hence in points where each m_i is either 1 or -1. Evidently these minima are disconnected, at odds with what happens when $\eta = 0$. Moreover there are 2^N such corners so one would expect the number of steady states to be exponentially large (in N). This is indeed the case. The difference between the MG with and without price impact can be summarized in the following table.

	$\eta = 0$	$\eta = 1$
Ergodicity	Preserved above α_c, broken below	broken
Nr of steady states	Unique above α_c, degenerate below	$\sim e^{\Sigma N}$, $\Sigma > 0$
Geometry of the steady state	Point-like above α_c, continuous below	disconnected
Dependence on Γ	Absent above α_c, present below	Absent
Agents' behaviour	Stochastic	Deterministic
Predictability	Present above α_c, absent below	present

6 The role of finite score memory: the case of Spherical Minority Games

Minority Games with finite score memories have attracted attention since they display a behaviour qualitatively different from the one of standard MGs [8]. In particular non-ergodic phases are absent. In such games agents 'forget' their past score valuations exponentially at a rate λ, and the score updates are performed according to the following (batch) rule

$$q_i(t+1) = (1-\lambda)q_i(t) - \sum_j J_{ij} s_j(t) - h_i, \tag{27}$$

where we have abbreviated $J_{ij} = \frac{2}{N}\sum_\mu \xi_i^\mu \xi_j^\mu$ and $h_i = \frac{2}{N}\sum_\mu \xi_i^\mu \Omega^\mu$, $\Omega^\mu = \sum_j \omega_j^\mu$. For $\lambda = 0$ one recovers the standard MG (no memory-loss). It is intuitively easy to understand that no non-ergodic behaviour and sensitivity to initial conditions is observed whenever $\lambda > 0$: as the outcomes of past rounds are taken into account with an exponentially vanishing weight as time progresses, the effects of the starting point decays as well, and initial conditions have no influence on the dynamics asymptotically. These properties make such model ideally suited for studying the onset of non-ergodicity by varying λ from 0 upwards. (It is worth adding that these games generate highly non-trivial and realistic time-series with bursts of fluctuations on as yet not understood time-scales [8, 24].) Even though such games are purely ergodic, they pose a serious challenge from the theoretical point of view (recall that in their ergodic regimes standard MGs are usually soluble). Due to the finiteness of the memory, there are no frozen agents when $\lambda > 0$. Unfortunately, all known theoretical approaches rely on the existence of the latter, so that further analytical understanding is impeded in the case of finite memory and absence of frozen agents.

Progress can here be made by considering so-called 'spherical' limits of the MG, addressed in [9, 25], inspired by work done in the physics of spin systems [26]. The hard variables $s_i(t) = \text{sgn}[q_i(t)] \in \{-1, +1\}$ are here replaced by 'soft' ones, $\phi_i(t) \in \mathbb{R}$, which in turn obey a spherical constraint $\sum_{i=1}^{N} \phi_i(t)^2 = N$. In a batch setup the update rules of the finite-memory spherical MG then read

$$q_i(t+1) = (1-\lambda)q_i(t) - \sum_j J_{ij}\phi_j(t) - h_i. \tag{28}$$

Note that $s_j(t)$ in (27) has been replaced $\phi_j(t)$. The spherical constraint is implemented through

$$\phi_i(t) = \frac{q_i(t)}{\rho(t)}, \tag{29}$$

where the multiplier $\rho(t)$ is chosen to ensure that $N^{-1}\sum_i \phi_i(t)^2 = 1$ at all times t. This model describes agents who use linear combinations of their two strategies.

A full theoretical analysis can then be performed in the ergodic phase. We do not give details here, but refer to [9, 25] for analyses of similar cases. Let it suffice to say that the volatility can be computed exactly without making any approximations at any stage[5].

Fig. 4 shows that the model without memory loss ($\lambda = 0$) shows ergodic behavior for $\alpha > \alpha_c$, and a non-ergodic phase with several branches of the volatility below the transition as described in detail in [25]. Simulations suggest that introducing finite memory $\lambda > 0$ removes the transition[6] and renders the model fully ergodic for all α. Still, unlike the non-spherical version, exact expressions can be found for the volatility, as indicated by the solid lines in Fig. 4.

[5]Expressions obtained for σ^2 in non-spherical MGs are typically only of an approximate nature, as certain correlations between agents or, equivalently, transient parts of the dynamics are neglected.

[6]Spherical models typically exhibit two types of ergodic phases, a frozen and an oscillatory one [9,25]. This appears to be the case for small $\lambda > 0$ as well, where as for large enough λ only the oscillatory phase appears to be present. The lower phase boundary of the frozen ergodic phase here warrants further investigation to fully confirm the absence of ergodicity breaking in the spherical model with finite memory.

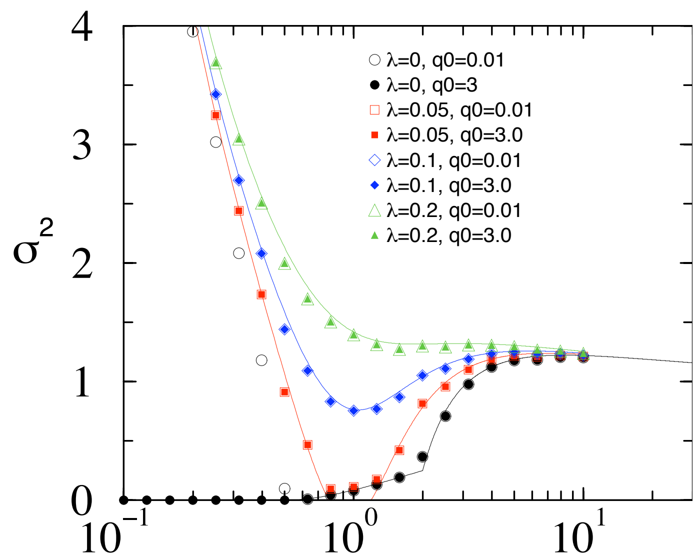

Figure 4: Volatility for the spherical MG with finite memory. Symbols: simulations are batch $N = 300$, 500 steps, 10 samples. Lines: theory.

7 Summary

Non-ergodic phases in Minority Games are not well understood. The ability of MGs to reproduce the empirical phenomenology (volatility clustering, fat-tailed price returns, etc) is more striking around the critical point separating the two phases [2,3, 6]. This suggests that markets operate on the border between efficient (non-ergodic) and inefficient (ergodic) phases. A deeper understanding of the former would thus be important not only from a purely theoretical viewpoint (from which obtaining solutions in the non-ergodic regime can be seen as an interesting mathematical challenge) but also from a practical side, e.g. to develop trading platforms based on the MG: phases with high volatility typically have tinier minorities, which make larger profits, than phases with low volatility. A reasonable mechanism for the switching between different phases might be based on the fact that more predictable markets would attract more investors hence decrease α (the ratio of the number information patterns divided by the number of agents), while more unpredictable markets would repel investors hence increase α. Thus α would naturally oscillate around the critical point, and potentially converge to result in an analogue of self-organised criticality. Interestingly, no model has been proposed so far that displays this dynamical feedback.

Acknowledgements

This work was supported by EU NEST No. 516446 COMPLEXMARKETS. The authors would like to thank all colleagues with whom they have been collaborating on the topic of Minority Games, and in particular Matteo Marsili and David Sherrington.

References

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

[2] Galla T, Mosetti G, Zhang Y-C 2006 *Anomalous fluctuations in Minority Games and related multi-agent models of financial markets*, review to appear in 'Anomalous fluctuation phenomena in complex systems' (Eds C. Riccardi and H.E. Roman)

[3] Challet D, Marsili M and Zhang Y-C 2005 *Minority Games* (Oxf ord University Press, Oxford UK)

[4] Coolen A C C 2005 *The Mathematical Theory of Minority Games* (Oxford University Press, Oxford UK)

[5] Moro E 2004 (to appear in *Advances in Condensed Matter and Statistical Mechanics*, ed. E. Korutcheva and R. Cuerno), preprint `cond-mat/0402651`

[6] De Martino A and Marsili M 2006 *J. Phys. A* **39** R465

[7] `http://www.unifr.ch/econophysics/minority`

[8] Challet D, De Martino A, Marsili M, Perez Castillo I 2006 *J. Stat. Mech.* P03004

[9] Galla T, Coolen ACC, Sherrington D 2003 *J. Phys. A: Math. Gen.* **36** 11159

[10] Cavagna A, Garrahan J P, Giardina I, and Sherrington D 1999 *Phys. Rev. Lett.* **83** 4429

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

[12] Savit R, Manuca R and Riolo R 1999 *Phys. Rev. Lett.* **82** 2203

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

[14] Challet D, Marsili M and Zhang Y-C 2000 *Physica A* **276** 284

[15] Garrahan J P, Moro E and Sherrington D 2000 *Phys. Rev E* **62** R9

[16] Cavagna A 1999 *Phys. Rev. E* **59** R3783

[17] Coolen ACC and Shayeghi N 2006 *J. Phys. A* **39** 13921

[18] De Martino A, Perez Castillo I and Sherrington D 2007 JSTAT P01006

[19] Heimel J A F and Coolen A C C 2001 *Phys. Rev. E* **63** 056121

[20] Galla T, Sherrington D 2005 *Eur. Phys. J. B* **46** 153

[21] Marsili M and Challet D 2001 *Phys. Rev. E* **64** 056138

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

[23] Heimel J A F, De Martino A 2001 *J.Phys. A: Math. Gen.* **34** L539

[24] Challet D, De Martino A and Marsili M, forthcoming.

[25] Galla T, Sherrington D 2005 *J. Stat. Mech.* P10009

[26] Berlin T H, Kac M 1952 *Physical Review* **86** 821