

Minority games with finite score memory

Damin Challet, Andrea De Martino, Matteo Marsili and
Isaac Castillo

Minority games with finite score memory

Damin Challet¹, Andrea De Martino², Matteo Marsili³ and Isaac Perez Castillo⁴

¹*Nomura Centre for Quantitative Finance, Mathematical Institute,
Oxford University, 24–29 St Giles’, Oxford OX1 3LB, United Kingdom*

²*INFN-SMC and Dipartimento di Fisica, Università di Roma “La Sapienza”, P.le A. Moro 2, 00185 Roma, Italy*

³*The Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11, 34100 Trieste, Italy*

⁴*Instituut voor Theoretische Fysica, Katholieke Universiteit Leuven, B-3001 Leuven, Belgium*

We analyze grand-canonical minority games with infinite and finite score memory and different updating timescales (from ‘on-line’ games to ‘batch’ games) in detail with various complementary methods, both analytical and numerical. We focus on the emergence of ‘stylized facts’ and on the production of exploitable information, as well as on the dynamic behaviour of the models. We find that with finite score memory no agent can be frozen, and that all the current analytical methods fail to provide satisfactory explanation of the observed behaviours.

PACS numbers: 05.10.Gg, 89.65.Gh, 02.50.Le, 87.23.Ge

I. INTRODUCTION

Few realistic agent-based models of financial markets can be understood in depth. Among these, minority games (MGs) [1, 2] are perhaps the most studied at a fundamental physical level, especially through a systematic use of spin-glass techniques [3–7]. The inclusion of many important aspects of market dynamics in the standard MG setup, however, often leads to considerable technical difficulties and raises new challenges for statistical mechanics. A particularly important modification concerns the memory of agents. In the original game, the learning dynamics on which traders base their strategic decisions is such that they remember all their past payoffs irrespective of how far in time they occurred. It is however reasonable to think that real traders tend to base their choices only on the most recent events. This rather natural extension was originally introduced in [8] for a model in which agents play the MG strategically. For the purpose of modeling financial markets, the relevant situation is instead that of price-taking, or naïve, agents. In this case, it has been argued that finite score memory gives rise to a surprisingly rich dynamical phenomenology [9]. In this paper, we will analyze such models in detail.

There are of course several ways to introduce a finite memory in the MG. A conceptually simple one is to fix a time window M (the ‘score memory’) during which agents keep exact track of their scores [10, 11]. The main advantage is that the game becomes Markovian of order M . This situation can be handled numerically for reasonably small M . Here, we are interested in the case in which M is of the order of the number of traders N , which is supposed to be very large (ultimately, the limit $N \rightarrow \infty$ will be considered). For the sake of simplicity, we choose to implement the situation in which scores are exponentially damped in time, which requires only a minor modification of the original equations. Furthermore, we shall focus on the grand-canonical MG (GCMG) [12], which is known to produce market-like fluctuation phenomena and whose properties have been shown to be extremely sensitive to the introduction of a finite memory [9]. Both

the ‘on-line’ and ‘batch’ versions of the model will be addressed. The two situations differ by the timescales over which agents update their status: in the former, the updating takes place at every time step; in the latter, it occurs roughly once every P time steps, with $P = \mathcal{O}(N)$. In particular, following [13], we use a parameter that interpolates between the ‘on-line’ and ‘batch’ models and study how the resulting fluctuation phenomena are affected by changes of updating timescales.

We shall proceed by defining the models (Sec. II) and exploring, via computer experiments, their behaviour (Sec. III). We will focus especially on the emergence and parameter-dependence of ‘stylized facts’, peculiar statistical regularities that are empirically observed in financial markets. In Sections IV and V we will characterize analytically the stationary states of the model with infinite and finite memory, respectively. We use static replica-based minimization techniques in order to compute the properties of the on-line game with infinite score memory. For the batch game, one can resort to dynamical mean-field theory. The resulting theory is exactly solvable in the case of infinite memory. For finite-memory games, we have to resort a Monte Carlo scheme, known as Eissfeller-Opper method [14], to extract the steady state from the dynamical mean-field equations. Finally, we formulate our concluding remarks in Sec. VI. For the sake of completeness, we add an Appendix in which the effects of finite score memory in the standard MG are discussed.

II. DEFINITION OF THE MODEL

In a minority game, at each time step t , N traders are faced with two choices: to buy or to sell; those who happen to be in the minority win. The game is repeated, and the traders’ actions are determined by a simple re-inforcement learning. Each agent i has his own fixed trading strategy $\mathbf{a}_i = \{a_i^\mu\}$ that specifies an action $a_i^\mu \in \{-1, 1\}$ for each of the $\mu = 1, \dots, P$ states of the world. Each component of every strategy is ran-

domly drawn from $\{-1, 1\}$ with uniform probability before the beginning of the game. The adaptation abilities of agent i are limited to choosing whether to participate or to withdraw from the market, denoted respectively by $n_i(t) = 1$ and $n_i(t) = 0$. At time t , the state of the world $\mu(t)$ is drawn equiprobably from $\{1, \dots, P\}$. Agent i sets his $n_i(t)$ according to the sign of his strategy score, denoted by $U_i(t)$. In particular, he plays $n_i(t)a_i^{\mu(t)}$ where $n_i(t) = \Theta[U_i(t)]$ and $\Theta(x)$ is the Heaviside step function. The total excess demand $A(t)$ at time t , namely the numerical difference between buyers and sellers, is

$$A(t) = \sum_{i=1}^N n_i(t)a_i^{\mu(t)}. \quad (1)$$

All the agents then update their score according to

$$y_i(t+1) = \left(1 - \frac{\lambda_i}{P}\right) y_i(t) - \frac{1}{P} a_i^{\mu(t)} A(t) - \frac{\epsilon_i}{P} \quad (2)$$

with $\lambda_i \geq 0$ a constant. The $(1 - \lambda_i/P)$ term is responsible for finite score memory: if $\lambda_i = 0$, the agent has infinite score memory, as usually assumed in MGs, while for $\lambda_i > 0$ he has an exponentially damped memory of his past scores. More precisely, the number of time steps needed to forget a fraction f_i of his/her payoff is $\ln(1 - f_i)/\ln(1 - \lambda_i/P)$; for $\lambda_i/P \ll 1$, it is proportional to P/λ_i . The payoff $-a_i^{\mu(t)}A(t)$ is that of a Minority Game. The last term ϵ_i sets a benchmark that agent i has to beat in order to participate in the market. For instance, ϵ_i can be thought of as the interest of a risk-free account (see [12, 15] for more details).

Although our analysis can be extended easily to a more heterogeneous case, we assume, for the sake of simplicity, that $\lambda_i = \lambda$ and that there are two groups of agents: those with $\epsilon_i = -\infty$, referred to as ‘producers’, who always take part in the market, and the rest, who have $\epsilon_i = \epsilon$ finite and are called ‘speculators’ [16]. Traders with $\epsilon > 0$ (resp. $\epsilon < 0$) are risk-averse (resp. risk-prone), i.e., they have an incentive to stay out of (resp. enter) the market. We denote by N_s and N_p the number of speculators and producers, respectively. The producers, being deterministic, inject information into the market, that the speculators try to exploit. This setup, which defines an ecology of market participants, has been introduced in [15] as the simplest tractable interacting agents model able to reproduce the ‘stylized facts’ of financial markets [12]. In the statistical mechanics approach, one is interested in the limit of large systems, in which $P, N_s, N_p \rightarrow \infty$ keeping the reduced number of agents $n_s = \lim_{P \rightarrow \infty} N_s/P$ and $n_p = \lim_{P \rightarrow \infty} N_p/P$ fixed.

Up to now, agents update their variables $n_i(t)$ at each time step. It is natural in financial markets to assume that traders prefer not to change their strategy every time step in order to avoid over-reacting, and also because estimating the performance of a strategy needs some time. This can be approximated in our model by

allowing agents to change their variable n_i every T time steps [13]. For the sake of simplicity, we assume that agents perform the updates synchronously. If $T \gg P$, the score update between t and $t + T$ is essentially an average of score increases over all the states of the world. In this limit, defining t' as t/T , one can rewrite (2) as

$$y_i(t' + 1) = (1 - \lambda_i)y_i(t') - \sum_{j=1}^N J_{ij}n_j(t') - \alpha\epsilon_i \quad (3)$$

with $\alpha = P/N = (n_s + n_p)^{-1}$ and quenched random couplings J_{ij} given by

$$J_{ij} = \frac{1}{N} \sum_{\mu=1}^P a_i^{\mu} a_j^{\mu} \quad (4)$$

Because this dynamics is equivalent to enumerating all $\mu \in \{1, \dots, P\}$, we assumed that all the states of the world always occur between t' and $t' + 1$. The neural network literature refers to models with $T = 1$ as ‘on-line’, whereas models such as (3) are called ‘batch’. The parameter T allows the interpolation between the former and the latter.

Batch minority games were introduced in [17] and solved later exactly with generating functionals [5], an exact dynamical method. The stationary states of (2) and (3) are, strictly speaking, different even in the thermodynamic limit. This is because agents in batch games have a longer auto-correlation: for instance, denoting $\phi_i = \langle n_i \rangle$ the time average of $n_i(t)$ in the stationary state, one has $\langle n_i n_j \rangle = \phi_i \phi_j$ in on-line games [3, 7], but not in batch games [7].

III. STYLIZED FACTS

The connection between minority game’s outcome A and real prices comes from relating A to the ‘excess demand’, and linking the evolution of the price $p(t)$ to it via

$$\log p(t+1) = \log p(t) + \frac{A(t)}{L} \quad (5)$$

where L is a constant, the ‘liquidity’, that will be hereafter fixed to 1 [18, 19]. While the original MG revealed insightful relationships between price fluctuations and predictability, it fails to reproduce the empirically observed market-like behaviour, in particular the so-called ‘stylized facts’. As a consequence, GCMGs were introduced to better mimic market dynamics. They are able to produce stylized facts such as fat-tailed price return distributions $P(A) \propto A^{-\beta}$ with $\beta \simeq 3.5$ and volatility clustering $\langle A(t)^2 A(t+\tau)^2 \rangle \propto \tau^{-\gamma}$ with $\gamma \simeq 0.3$; these exponents are close to those measured in real markets. They are however unable to reproduce well documented over-diffusive price behavior [20, 21], because the MG induces a mean-reverting process.




FIG. 1: Return as a function of time for the same realisation of a GCMG with biased initial strategy scores and $\lambda = 0$ (black lines), and $\lambda = 0.003$ (red lines). System set in the critical window, $P = 25$, $n_s = 40$, $\epsilon = 0.01$

The presence of stylized facts in the GCMG with $\lambda = 0$ was linked to a too small signal-to-noise ratio, suggesting that marginal efficiency is a necessary condition for the existence of stylized facts [12]. In other words, in infinite systems, stylized facts only occur at a the phase transition, whereas in finite systems, critical-like phenomenology is observed in a critical window around the critical point, which shrinks as the system size increases (see [12] for more details and for a way of keeping alive stylized facts in infinite systems). Therefore, the stylized facts of this model cannot be studied by the methods of statistical mechanics.

While the *ability* of the GCMG to produce stylized facts was emphasised in previous work, Ref. [9] pointed out two delicate problems of GCMG regarding stylized facts. The first one is the dependence of stylized facts on initial conditions: assume that one given realization of the game produces stylized facts; changing slightly the initial conditions $y_i(0)$ is then enough to destroy them, leading to scenarios with only Gaussian price changes. The second problem is the following: for a given set of parameters, one realization of the game may produce stylized facts but another not. Both problems are due to the coincidence of fixed disorder in the strategies and infinite score keeping. This ultimately motivates the introduction of finite score memory. Fig. 1 compares $A(t)$ for the same game with $y_i(0) \neq 0$, once with $\lambda = 0$ and once with $\lambda > 0$. In the latter case, fluctuations have a characteristic pattern as a function of time: they first decrease to a very small value, stay at this level for a time interval of order $1/\lambda$ and then increase, producing fat-tailed price returns. The same kind of volatility behavior occurs in the original MG (see Appendix A).

However, finite score memory is a double-edged sword.




FIG. 2: Absolute-valued price return auto-correlation function for increasing λ for a given realisation of an on-line game; inset: same data in double-linear scale (10^7 iterations per run, after $200P$ iterations; $P = 20$, $n_s = 20$, $n_p = 1$, $\epsilon = 0.01$)

First of all, certain empirical stylized facts, in particular ‘volatility clustering’, suggest the presence of long-memory effects in markets. If one defines volatility clustering via the requirement

$$C(\tau) \equiv \frac{\langle |A(t)||A(t+\tau)| \rangle}{\langle A(t)^2 \rangle} \simeq \tau^{-\gamma} \quad (6)$$

it is easily understood that any positive value of λ destroys this power-law dependence, because of the cut-off that it imposes at times of the order of $P/|\ln(1-\lambda)|$. Fig. 2, by displaying $C(\tau)$ for $\tau \leq 1000$, is not able to show this effect for $\lambda < 0.001$; the cut-off is clear for larger λ . Not only $C(\tau)$ is cut off, but the loss of memory induces a negative $C(\tau)$ at large τ (see inset).

The other potentially nefarious effect of finite score memory is threaten the power-law tails: GCMGs pro-




FIG. 3: Average price return distribution function of on-line games for increasing λ . Average of 100 samples on 10^6 iterations per run, after $200P$ iterations; $P = 20$, $n_s = 20$, $n_p = 1$, $\epsilon = 0.01$.




FIG. 4: Average price return distribution function of on-line games for increasing λ . Average of 100 samples on 10^6 iterations per run, after $200P$ iterations; $P = 20$, $n_s = 40$, $n_p = 1$, $\epsilon = 0.01$.

duce power-law tailed A because of a volatility feed-back [12]. This feed-back needs some time to establish, hence if the memory length associated to λ is smaller than this time, power-law tails should disappear. This is indeed the case (see Figs 3 and 4), in a queer way: the central part of $P(A)$ is exponential, but the size of the support of $P(A)$ actually increases because of the appearance of two peaks. Note that maximum value λ^* of λ for which stylized facts are preserved depends on the system's parameter, as shown by these two figures.

We conclude that, in practice, a sufficiently small $\lambda \geq \lambda^* \simeq 0.001$ preserves the salient stylized facts: the noise of $C(\tau)$ for $\tau \geq 100$ in GCMG or in financial market data is such that the values of λ that preserve power-law tails of $P(A)$ do so for $C(1)$. Therefore, the introduction of finite score memory does not affect significantly the market-like phenomenology produced by GCMG. The value of λ^* obtained has to be contrasted with the small typical time-window used e.g. in [10, 11].

Let us now investigate the effect of updating the $n_i(t)$ every T time steps (see Fig. 5). Increasing T , one interpolates between on-line games ($T = 1$) and batch games ($T \gg P$). One clearly sees that this destroys large price changes: while $P(A)$ has power-law tails when $T = 1$, it is gradually transformed into an exponential distribution for $T \gg P$. This is because large price changes is associated to a given pattern μ [9, 11] that varies as a function of time, as illustrated by Fig. 6. However, it should be noted that $T \simeq P$ gives rise to a cleaner power-law.

Batch games display the same kind of transition when λ is increased. Fig. 7 shows that the largest value of λ such that $P_\lambda(A) \simeq P_0(A)$ is smaller and around 0.002 for $n_s = 20$, $P = 20$. Interestingly, the distribution of price changes $P(A)$ is more or less stable with respect to λ as long as $\lambda < 0.004$ in this figure. This is an important condition for the use of finite memory in these models.




FIG. 5: Average price return distribution function for increasing T . Average of 100 samples on 1,000,000 iterations per run, after $200P$ iterations; $P = 20$, $n_s = 40$, $n_p = 1$, $\epsilon = 0.01$.




FIG. 6: Price return as a function of time for various patterns μ . A symbol is plotted only when $\mu(t) = \mu$.

IV. STATIONARY STATE WITH $\lambda = 0$

This section characterizes the steady state of both the on-line (Eq. (2)) and batch GCMG (Eq. (3)) for $\lambda = 0$. The relevant macroscopic observables are as usual the 'volatility' σ^2 and the 'predictability' H , given respectively by

$$\sigma^2 = \frac{\langle A^2 \rangle}{P} \quad H = \frac{\overline{\langle A | \mu \rangle^2}}{P} \quad (7)$$

where $\langle \dots \rangle$ and $\langle \dots | \mu \rangle$ denote time averages in the stationary state, the latter conditioned on the occurrence of the piece of information μ , and the over-line denotes an average over information patterns. In on-line games, these two quantities are linked by

$$\sigma^2 = H + \phi - G \quad (8)$$




FIG. 7: Average price return distribution function of batch games for increasing λ . Average of 100 samples on 1,000,000 batch iterations per run, after $200P$ batch iterations; $P = 20$, $n_s = 20$, $n_p = 1$, $\epsilon = 0.01$.

where $G = \sum_{i=1}^{N_s} \phi_i^2 / P$, $\phi_i = \langle n_i \rangle$ denoting the probability that speculator i joins the market in the steady state. The normalizing P factors have been introduced in order to ensure that all quantities remain finite when P , N_s , $N_p \rightarrow \infty$. In addition to these, an important role is played by the relative number of active speculators

$$\phi = \frac{1}{N_s} \sum_{i=1}^{N_s} \phi_i \quad (9)$$

as well as by the number of active speculators per pattern, i.e. $n_{\text{act}} = n_s \phi$.

A. On-line GCMG: static approach

Partial results from replica-based calculus are reported in Ref. [12], and are based on the existence of a global quantity H_ϵ that is minimized by the dynamics. Finding H_ϵ relies on the prescription given in Ref. [22]: first derive the continuous-time stochastic differential version of Eq. (2), which reads

$$\dot{y}_i(\tau) = -\overline{a_i \langle n_i A \rangle} - \epsilon + \eta_i(\tau) \quad (10)$$

where $\tau = t/P$ is the intrinsic time of the GCMG, $\eta_i(t)$ is a zero-average Gaussian noise with $\langle \eta_i(\tau) \eta_j(\tau') \rangle = \frac{1}{N} \overline{a_i a_j \langle A^2 \rangle} \delta(\tau - \tau')$, and the over-line denotes an average over information patterns: $\overline{\dots} = (1/P) \sum_{\mu=1}^P \dots$. The deterministic term of this equation can be interpreted as the gradient of H_ϵ , so that $H_\epsilon = H_0 + 2\epsilon \sum_{i=1}^{N_s} \phi_i$. Stationary states correspond then to the minima of H_ϵ . When the stationary state is unique, its equilibrium properties are entirely determined by H_ϵ ; otherwise, one has to supplement H_ϵ with a self-consistent equation for computing correlations [22]. Regarding H_ϵ as a cost function,

one may compute the minima of H_ϵ from the partition function $Z(\beta) = \text{Tr} e^{-\beta H_\epsilon}$. The typical properties of the minimum of H_ϵ , i.e. of the $\beta \rightarrow \infty$ limit of free energy, require the evaluation of the quenched-disorder average $[\log Z]_{\text{dis}}$, which is performed via the replica trick $[\log Z]_{\text{dis}} = \lim_{n \rightarrow 0} \log [\langle Z^n \rangle]_{\text{dis}} / n$.

Reference [12] reported plots of the exact solution. Here we give the final results of the calculus only, as the latter is standard (see however [16] for more details). With replica symmetric ansatz, the free energy $f = (1/\beta \partial) \ln Z / \partial \beta$, which corresponds to H_ϵ in the limit $\beta \rightarrow \infty$, is given by

$$f(g, r) = \frac{\alpha}{2\beta} \log \left[1 + \frac{2\beta(G-g)}{\alpha} \right] + \frac{\rho+g}{1+\chi} + \frac{\alpha\beta}{2} (RG - rg) - \frac{1}{\beta} \langle \log \int_0^1 d\pi e^{-\beta V_z(\pi)} \rangle \quad (11)$$

where we found it convenient to define the ‘‘potential’’

$$V_z(\pi) = -\frac{\alpha\beta(R-r)}{2} \pi^2 - \sqrt{\alpha r} z \pi + 2\epsilon \pi \quad (12)$$

so that the last term of f looks like the free energy of a particle in the interval $[0, 1]$ with potential $V_z(\pi)$ where z plays the role of disorder. $G = \sum_{i=1}^{N_s} \phi_i^2$ is the self-overlap and g , R , r are Lagrange multipliers.

The four saddle point equations have exactly the same form as MG without the 0 strategy:

$$\frac{\partial f}{\partial g} = 0 \quad \Rightarrow \quad r = \frac{4(\rho+g)}{\alpha^2(1+\chi)^2} \quad (13)$$

$$\frac{\partial f}{\partial G} = 0 \quad \Rightarrow \quad \beta(R-r) = -\frac{2}{\alpha(1+\chi)} \quad (14)$$

$$\frac{\partial f}{\partial R} = 0 \quad \Rightarrow \quad G = \langle \langle \pi^2 \rangle \rangle_z \quad (15)$$

$$\frac{\partial f}{\partial r} = 0 \quad \Rightarrow \quad \beta(G-g) = \frac{\langle \langle \pi z \rangle \rangle_z}{\sqrt{\alpha r}} \quad (16)$$

In the limit $\beta \rightarrow 0$ we can look for a solution with $g \rightarrow G$ and $r \rightarrow R$. It is convenient to define

$$\chi = \frac{2\beta(G-g)}{\alpha}, \quad \text{and} \quad \zeta = -\sqrt{\frac{\alpha}{r}} \beta(R-r) \quad (17)$$

and to require that they stay finite in the limit $\beta \rightarrow \infty$. The potential can then be rewritten as

$$V_z(\pi) = \sqrt{\alpha r} \left[\zeta \frac{\pi^2}{2} - \pi \left(z - \frac{2\epsilon}{\sqrt{\alpha r}} \right) \right] \quad (18)$$

The averages are easily evaluated since, in this case, they are dominated by the minimum of the potential $V_z(\pi)$. Let K be $\epsilon(1+\chi)$, the minimum of $V_z(\pi)$ is at $\pi = 0$ for $z \leq \zeta K$ and at $\pi = +1$ for $z \geq \zeta(1+K)$. For $\zeta K < z < \zeta(1+K)$, the minimum is at $\pi = z/\zeta - K$. With this we find,

$$\langle \langle \pi z \rangle \rangle = \frac{1}{2\zeta} \left\{ \text{erf}[(1+K)\zeta/\sqrt{2}] - \text{erf}(K\zeta/\sqrt{2}) \right\} \quad (19)$$

and

$$\begin{aligned} \langle\langle \pi^2 \rangle\rangle &= G = \frac{1}{\zeta\sqrt{2\pi}} \left[(K-1)e^{-(1+K)^2\zeta^2/2} - Ke^{-K^2\zeta^2/2} \right] \\ &+ \frac{1}{2} \left(K^2 + \frac{1}{\zeta^2} \right) \left(\operatorname{erf}[(1+K)\zeta/\sqrt{2}] - \operatorname{erf}(K\zeta/\sqrt{2}) \right) \\ &+ \frac{1}{2} \operatorname{erfc}[(1+K)\zeta/\sqrt{2}] \end{aligned} \quad (20)$$

The fraction of agents who never enter into the market is then $\phi_0 = (1 + \operatorname{erf}(\zeta K/\sqrt{2}))/2$, and the ones who always participate is $\phi_1 = \operatorname{erfc}[(1+K)\zeta/\sqrt{2}]/2$. The pdf of the π_i is given by

$$\mathcal{P}(\pi) = \phi_0 \delta(\pi) + \phi_1 \delta(\pi - 1) + \frac{\zeta}{\sqrt{2\pi}} e^{-(\zeta(\pi+K))^2/2} \quad (21)$$

And the average number of agents in the market $\phi = \langle \pi \rangle$ where the average is over $\mathcal{P}(\pi)$, is

$$\begin{aligned} \phi &= \phi_1 + \frac{1}{\zeta\sqrt{2\pi}} \left(e^{-K^2\zeta^2/2} - e^{-\zeta^2(1+K)^2/2} \right) \\ &+ \frac{K}{2} \left(\operatorname{erf}(K\zeta/\sqrt{2}) - \operatorname{erf}[(1+K)\zeta/\sqrt{2}] \right) \end{aligned}$$

Observing that $\zeta = \sqrt{\alpha/(\rho+G)}$, one finally finds that ζ is fixed as a function of α and ρ by the equation

$$\begin{aligned} \frac{\alpha}{\zeta^2} &= \rho + \frac{1}{\zeta\sqrt{2\pi}} \left[(K-1)e^{-(1+K)^2\zeta^2/2} - Ke^{-K^2\zeta^2/2} \right] \\ &+ \frac{1}{2} \left(K^2 + \frac{1}{\zeta^2} \right) \left(\operatorname{erf}[(1+K)\zeta/\sqrt{2}] - \operatorname{erf}(K\zeta/\sqrt{2}) \right) \\ &+ \frac{1}{2} \operatorname{erfc}[(1+K)\zeta/\sqrt{2}] \end{aligned} \quad (22)$$

which has to be solved numerically. With some more algebra, one easily finds :

$$K = \epsilon \left[1 - \frac{\operatorname{erf}[(1+K)\zeta/\sqrt{2}] - \operatorname{erf}(K\zeta/\sqrt{2})}{2\alpha} \right]^{-1} \quad (23)$$

These two last equations form a close non-linear set of equations.

The calculus gives finally $H_\epsilon = \lim_{\beta \rightarrow \infty} f$

$$H_\epsilon = \frac{n_p + n_s G(\zeta, K)}{(1 + \chi)^2} + 2\epsilon \phi(K, \zeta) \quad (24)$$

The stationary is unique if $H_\epsilon \neq 0$, which is the case as long as $\epsilon \neq 0$ and for $n_s \leq n_s^*(n_p, \epsilon)$ if $\epsilon = 0$.

The fluctuations are given by

$$\sigma^2 = \epsilon^2 \frac{n_p + n_s G}{K^2} + n_s(\phi - G). \quad (25)$$

One can also show that $H_\epsilon \propto \epsilon^2$ for small ϵ : according to Eq. (24), this hold if $K \rightarrow 0 < K_0 < \infty$, which can be seen by expanding Eq. (23) in powers of ϵ .

B. Batch GCMG: dynamical approach

We consider the batch-GCMG dynamics, which we recast as

$$y_i(t+1) = y_i(t) - \sum_{j=1}^N J_{ij} n_j(t) - \alpha \epsilon_i + h_i(t) \quad (26)$$

with

$$\epsilon_i = \begin{cases} \epsilon & \text{for } 1 \leq i \leq N_s \\ -\infty & \text{for } N_s + 1 \leq i \leq N_s + N_p \equiv N \end{cases}$$

The subscripts s and p denote speculators and producers, respectively. The relevant dynamical variable is $n_i(t) = \Theta[y_i(t)]$ (for producers, $n_i(t) = 1$), while the random couplings J_{ij} are given by (4), $J_{ij} = (1/N) \sum_\mu a_i^\mu a_j^\mu$, with $a_i^\mu \in \{-1, 1\}$ iid quenched random variables with uniform probability distribution. For simplicity, we set

$$\alpha = \frac{P}{N} = \frac{1}{n_s + n_p} \quad (27)$$

with $n_s = N_s/P$ and $n_p = N_p/P$. The external sources $h_i(t)$ have been added in order to probe the system against small perturbations. In the following, we denote averages over all possible time-evolutions (paths), i.e. realizations of (26), by double brackets $\langle\langle \dots \rangle\rangle$.

The standard tool for investigating the dynamics of statistical systems with quenched disorder is the path-integral method à la De Dominicis, based on the evaluation of the generating functional

$$Z[\psi] = \left[\left\langle \left\langle \exp \left[-i \sum_{i,t} n_i(t) \psi_i(t) \right] \right\rangle \right\rangle_{\text{dis}} \right] \quad (28)$$

from which disorder-averaged site-dependent correlation functions of all orders can be derived via such identities as

$$\langle\langle n_i(t) \rangle\rangle_{\text{dis}} = i \lim_{\psi \rightarrow 0} \frac{\partial Z[\psi]}{\partial \psi_i(t)} \quad (29)$$

$$\langle\langle n_i(t) n_j(t') \rangle\rangle_{\text{dis}} = - \lim_{\psi \rightarrow 0} \frac{\partial^2 Z[\psi]}{\partial \psi_i(t) \partial \psi_j(t')} \quad (30)$$

In turn, macroscopic (auto)-correlation and response functions like

$$C(t, t') = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \langle\langle n_i(t) n_i(t') \rangle\rangle_{\text{dis}} \quad (31)$$

$$G(t, t') = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{\partial \langle\langle n_i(t) \rangle\rangle_{\text{dis}}}{\partial h_i(t')} \quad (32)$$

can in principle be evaluated by simply taking derivatives of Z with respect to the sources $\{\psi_i, h_i\}$. The calculation of Z in the limit $N \rightarrow \infty$ leads, via a saddle-point integration, to the identification of a non-Markovian single

(‘effective’) agent process that provides a complete description of the original Markovian multi-agent process (26). This procedure requires a straightforward variation of the cases dealt with in the existing literature and we will not report it in detail. For our purposes, it will suffice to say that the effective dynamics for speculators is given by

$$y(t+1) = y(t) - \alpha \sum_{t' \leq t} (I + G)^{-1}(t, t') n(t') - \alpha \epsilon + h(t) + \sqrt{\alpha} z(t) \quad (33)$$

with $n(t) = \Theta[y(t)]$, whereas $n(t) = 1$ for producers always. Here, I is the identity matrix, G is the response function (32), and $z(t)$ is a Gaussian noise with zero average and time correlations

$$\langle z(t) z(t') \rangle = [(I + G)^{-1} C (I + G^T)^{-1}](t, t') \quad (34)$$

with C the correlation function (31). Notice that the coupling between the two groups is provided in essence by the noise term, since both speculators and producers contribute to C . In fact, the correlation function can be written as

$$C(t, t') = n_s \alpha C_s(t, t') + n_p \alpha \quad (35)$$

where $n_s \alpha$ (resp. $n_p \alpha$) is the fraction of speculators (resp. producers), and C_s (resp. C_p) denote the correlation function of speculators (resp. producers). For the response function we have, similarly,

$$\begin{aligned} G(t, t') &= n_s \alpha G_s(t, t') + n_p \alpha G_p(t, t') \\ &= n_s \alpha G_s(t, t') \end{aligned} \quad (36)$$

producers being ‘frozen’ at $n = 1$ and thus insensitive to small perturbations.

Assuming time-translation invariance,

$$\lim_{t \rightarrow \infty} C(t + \tau, t) = C(\tau) \quad (37)$$

$$\lim_{t \rightarrow \infty} G(t + \tau, t) = G(\tau) \quad (38)$$

finite integrated response,

$$\lim_{t \rightarrow \infty} \sum_{t' \leq t} G(t, t') < \infty \quad (39)$$

and weak long-term memory,

$$\lim_{t \rightarrow \infty} G(t, t') = 0 \quad \forall t' \text{ finite} \quad (40)$$

ergodic steady states of (33) can be characterized in terms of a couple of order parameters, namely the persistent autocorrelation

$$c = \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \sum_{t < \tau} C(\tau) \quad (41)$$

and the integrated response (or susceptibility)

$$\chi = \lim_{\tau \rightarrow \infty} \sum_{t < \tau} G(\tau) \quad (42)$$

From (35) and (36) we get

$$c = n_s \alpha c_s + n_p \alpha \quad (43)$$

$$\chi = n_s \alpha \chi_s \quad (44)$$

where c_s and χ_s are the persistent autocorrelation and susceptibility of speculators. For $\lambda = 0$ one can formulate, inspired by computer experiments, a simple Ansatz for the dynamics of scores that allows to calculate these quantities exactly as functions of c and χ , so that from (43) and (44) one may retrieve the values of the persistent order parameters for any n_s and n_p . In this case, ergodicity breaks down as χ diverges at certain critical values of the parameters, thus violating (39). This analysis, including the dynamical phase transition, reproduces the phenomenology of the batch-GCMG without memory remarkably well, at least in the ergodic regime.

The score $y(t)$ of speculators either grows linearly with time as $y(t) \simeq v t$ (in which case the agent is ‘frozen’ at inactivity with $n(t) = 0$ if $v < 0$ or activity with $n(t) = 1$ if $v > 0$), or keeps oscillating about $y(t) = 0$ (in which case the agent is ‘fickle’) [23]. In order to distinguish between the two situations, we introduce the variable $\tilde{y}(t) = y(t)/t$. Using this, we can sum (33) over time to obtain

$$\begin{aligned} \tilde{y}(t+1) - \frac{1}{t} y(1) &= -\frac{\alpha}{t} \sum_{t', t''} (I + G)^{-1}(t', t'') n(t'') \\ &\quad - \alpha \epsilon + \frac{\sqrt{\alpha}}{t} \sum_{t'} z(t') \end{aligned} \quad (45)$$

where $n(t) = \Theta[\tilde{y}(t)]$. In the limit $t \rightarrow \infty$, the above leads, via (37–40), to a simple equation for the quantity $\tilde{y} = \lim_{t \rightarrow \infty} \tilde{y}(t)$:

$$\tilde{y} = -\frac{\alpha n}{1 + \chi} - \alpha \epsilon + \sqrt{\alpha} z \quad (46)$$

where χ is given by (42), $n = \lim_{\tau \rightarrow \infty} (1/\tau) \sum_{t \leq \tau} n(t)$, and $z = \lim_{\tau \rightarrow \infty} (1/\tau) \sum_{t \leq \tau} z(t)$ is a zero-average Gaussian rv with variance

$$\langle z^2 \rangle = \frac{1}{\tau \tau'} \sum_{t \leq \tau, t' \leq \tau'} \langle z(t) z(t') \rangle = \frac{c}{(1 + \chi)^2} \quad (47)$$

Defining $\gamma = \sqrt{\alpha}/(1 + \chi)$, we can proceed as usual by separating the frozen speculators from the fickle ones. We have the following situation: for $\tilde{y} > 0$ the effective speculator is always active ($n = 1$) and $z > \gamma + \sqrt{\alpha} \epsilon$; for $\tilde{y} < 0$ the effective speculator is always inactive ($n = 0$) and $z < \sqrt{\alpha} \epsilon$; for $\tilde{y} = 0$ the effective speculator is fickle, $n = \frac{z - \sqrt{\alpha} \epsilon}{\gamma}$ and $\sqrt{\alpha} \epsilon < z < \gamma + \sqrt{\alpha} \epsilon$. So we have

$$\begin{aligned} c_s &= \langle \Theta(z - \gamma - \sqrt{\alpha} \epsilon) \rangle \\ &\quad + \langle \Theta(z - \sqrt{\alpha} \epsilon) \Theta(\gamma + \sqrt{\alpha} \epsilon - z) \left(\frac{z - \sqrt{\alpha} \epsilon}{\gamma} \right)^2 \rangle \end{aligned} \quad (48)$$

with brackets denoting an average over z . For χ_s , we use the fact that the noise $z(t)$ formally acts like an external source in (46), so that $\chi_s = (1/\sqrt{\alpha})\langle \frac{\partial n}{\partial z} \rangle$. This gives

$$\chi_s = \frac{1}{\gamma\sqrt{\alpha}} \langle \Theta(z - \sqrt{\alpha}\epsilon)\Theta(\gamma + \sqrt{\alpha}\epsilon - z) \rangle \quad (49)$$

One can also calculate the average activity level of speculators as

$$\begin{aligned} \phi \equiv \langle n \rangle &= \langle \Theta(z - \gamma - \sqrt{\alpha}\epsilon) \rangle + \\ &\langle \Theta(z - \sqrt{\alpha}\epsilon)\Theta(\gamma + \sqrt{\alpha}\epsilon - z) \left(\frac{z - \sqrt{\alpha}\epsilon}{\gamma} \right) \rangle \end{aligned} \quad (50)$$

from which the number of active speculators per pattern follows as $n_{\text{act}} = n_s\phi$, and the fraction of frozen speculators as

$$f = \langle \Theta(z - \gamma - \sqrt{\alpha}\epsilon) \rangle + \langle \Theta(\alpha\epsilon - z) \rangle = f_1 + f_0 \quad (51)$$

where f_1 (resp. f_0) stands for the fraction of always active (resp. inactive) speculators. Inserting (48) and (49) into (43) and (44) one obtains two equations that can be solved self-consistently for c and χ . These equations can be analyzed for any ϵ . For the sake of simplicity, we focus on the case $\epsilon = 0$, in which the averages over z take a particularly simple form. We have

$$\begin{aligned} c = n_p\alpha + n_s\alpha &\left[\frac{1}{2} \left(1 - \text{erf} \sqrt{\frac{\alpha}{2c}} \right) \right. \\ &\left. + \frac{c}{2\alpha} \text{erf} \sqrt{\frac{\alpha}{2c}} - e^{-\frac{\alpha}{2c}} \sqrt{\frac{c}{2\pi\alpha}} \right] \end{aligned} \quad (52)$$

$$\frac{\chi}{1+\chi} = \frac{n_s}{2} \text{erf} \sqrt{\frac{\alpha}{2c}} \quad (53)$$

whereas ϕ reads

$$\phi = \frac{1}{2} \left(1 - \text{erf} \sqrt{\frac{\alpha}{2c}} \right) + \sqrt{\frac{c}{2\pi\alpha}} (1 - e^{-\frac{\alpha}{2c}}) \quad (54)$$

As a quick consistency check, one can see, starting from (53) and with minor manipulations, that a divergence of the susceptibility (i.e. the violation of (39) with consequent ergodicity breaking) for $n_p = 1$ occurs at $n_s = n_s^* = 2/\text{erf}(\xi^*)$ where ξ^* is the solution of the transcendental equation $e^{-\xi^2} = \xi\sqrt{\pi}$. The result, $n_s^* = 4.14542\dots$, is in full agreement with the replica results of both the on-line and the batch-model.

In order to compare with the computer experiments discussed above, we analyze the fraction of active speculators and the volatility σ^2 obtained from (52-53) at $n_p = 1$. As said above, the former is just $n_{\text{act}} = n_s\phi$. As for the latter, it is formally given by

$$\sigma^2 = \lim_{t \rightarrow \infty} \langle z(t)z(t) \rangle / \alpha \quad (55)$$




FIG. 8: GCMG with $\lambda = 0$: number of active speculators per pattern $n_{\text{act}} = n_s\phi$ (top) and volatility (bottom) at $n_p = 1$ as a function of n_s for different values of ϵ . The dashed vertical line marks the position of the critical point n_s^* . Markers denote results from computer experiments.

It is possible to derive an approximate expression for the above limit in terms of persistent order parameters assuming that the retarded self-interaction of fickle speculators is negligible, that is, by neglecting the agent's auto-correlation. This leads to [5]

$$\alpha\sigma^2 = \frac{n_p + n_s f_1}{(1+\chi)^2} + \frac{1}{4}n_s(1-f). \quad (56)$$

Comparing with the replica result, one sees that σ^2 comprises a H_0 term, as expected from the relationship (25). However, the other terms differ, as in the standard MG. In Fig. 8 we report the behavior of ϕ and σ^2 obtained from (52-53) at $n_p = 1$. Dynamical results for the batch model are in excellent agreement with the simulations of the batch-GCMG and reproduce qualitatively the phenomenology of the on-line-GCMG. For the sake of completeness, we also report (see Fig. 9), for $\epsilon = 0$, the critical line $n_s^*(n_p)$ where χ diverges and ergodicity breaks down. After some algebra, it turns out to be given




FIG. 9: GCMG with $\lambda = 0$: critical line n_s^* vs n_p , where χ diverges and (39) is violated at $\epsilon = 0$. The dynamics is ergodic for $n < n_s^*$.




FIG. 10: ϕ_1 and ϕ_0 as a function of time for games with $\lambda = 0$ (black lines), 0.01 (red lines) and 0.1 (green lines). $P = 100$, $n_s = 4$, $\epsilon = 0.05$

by $n_s^*(n_p) = 2/\text{erf}(\xi^*)$, where $\xi^* \equiv \xi^*(n_p)$ is the solution of

$$\frac{e^{-\xi^2}}{\xi\sqrt{\pi}} = (n_p - 1)\text{erf}(\xi) + 1 \quad (57)$$

V. STATIONARY STATE WITH $\lambda > 0$

One may expect that the stationary state of Minority Games with $\lambda > 0$ depends smoothly on λ in ergodic regions, and indeed some important quantities such as the fluctuations and H do behave this way. More surprising is the vanishing of frozen agents: Fig. 10 reports that both ϕ_0 and ϕ_1 seem to cancel for any positive λ , although this may not appear for finite-time simulations at small values of λ ; the measure was done from $t = 1000$ and counts the fraction of agents that are never out of the market and in the market, respectively.

A. On-line GCMG: static approach

In many versions of the minority game with infinite score memory ($\lambda = 0$) and naïve agents, a phase transition takes place. H behaves like a physical order parameter which is minimized by the dynamics. The latter is ergodic as long as $H > 0$. However, when $H = 0$, the stationary state is not unique, and the dynamics becomes non-ergodic: the stationary state is selected by the initial score valuation $y_i(0)$; this happens in particular in the original MG and in the present GCMG with $\epsilon = 0$ and $\lambda = 0$. On the other hand, it is obvious that if $y_i(0)$ is gradually forgotten, the stationary state cannot depend anymore on $U_i(0)$. This is precisely what the introduction of the finite score memory does: when

$\lambda > 0$, the dynamics is ergodic, and accordingly the stationary state is unique. This would be compatible with a minimized quantity that is not cancelled anymore by the dynamics; intuitively, this means that a new term is added to H_ϵ . Unfortunately, finding such function turns out to be impossible in this case, which rules out the use of replica-based approach.

B. Batch GCMG: dynamical approach

Let us now turn to the effective process (33) with finite score memory. As discussed above, in this case scores do not diverge with time, i.e. $\lim_{t \rightarrow \infty} y(t) < \infty$, and it is no longer possible to separate frozen agents from fickle ones by the use of the quantity $\tilde{y} = \lim_{t \rightarrow \infty} y(t)/t$. Indeed, proceeding as done for the case $\lambda = 0$, one obtains, in place of (46), the condition

$$\lambda \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{t' \leq t} y(t') = -\frac{\alpha\phi}{1+\chi} - \alpha\epsilon + \sqrt{\alpha}z \quad (58)$$

where the term on the l.h.s. is finite. In addition, the fact that the fraction of frozen players undergoes an extremely slow dynamics ultimately suggesting that all agents are fickle in the stationary state indicates the necessity of a different analytical approach when $\lambda \neq 0$. Unfortunately, we have been unable to capture the peculiarities of the score dynamics with a simple Ansatz.

Some general hints can be obtained by calculating explicitly the first time step of (33). For simplicity, we henceforth adopt the shorthand $\langle z(t)z(t') \rangle = L(t, t')$ for (34). Furthermore, we assume an initial condition $y(0)$ for (33) ensuring that all speculators are active at time 0, so that $C(0, 0) = 1$ and (by causality) $G(0, 0) = 0$ and $L(0, 0) = 1$. The transition probability to pass from $y(0)$ to $y(1)$ is given by

$$p[y(1)|y(0)] = \frac{1}{\sqrt{\alpha}} e^{-\frac{1}{2\alpha}[y(1)-(1-\lambda)y(0)-\Theta(0)+\alpha n(0)+\alpha\epsilon]^2} \quad (59)$$

As a consequence, we have

$$\begin{aligned} C(1, 0) &= \int dy(1)dy(0)p[y(1)|y(0)]p(y(0))n(1)n(0) \\ &= \frac{1}{2} \left[1 - \text{erf} \left(\frac{\alpha n(0) + \alpha\epsilon - h(0) - (1-\lambda)y(0)}{\sqrt{2\alpha}} \right) \right] \\ &= \frac{1}{2} \left[1 - \text{erf} \left(\frac{\alpha + \alpha\epsilon - (1-\lambda)y(0)}{\sqrt{2\alpha}} \right) \right] \quad (60) \end{aligned}$$

where we set $h(0) = 0$ and used the fact that $n(0) = 1$, and

$$\begin{aligned} G(1, 0) &= \frac{\partial}{\partial h(0)} \int dy(1)dy(0)p[y(1)|y(0)]p(y(0))n(1) \\ &= \frac{1}{\sqrt{2\pi\alpha}} \exp \left\{ -\frac{1}{2\alpha} [\alpha + \alpha\epsilon - (1-\lambda)y(0)]^2 \right\} \quad (61) \end{aligned}$$

From these, we can evaluate (recalling that in the steady state (55) holds) the time-dependent volatilities

$$L(1,0) = \sum_{t,t'} (I+G)^{-1}(1,t)C(t,t')(I+G^T)^{-1}(t',0) \\ = -G(1,0) + C(1,0) \quad (62)$$

and

$$L(1,1) = \sum_{t,t'} (I+G)^{-1}(1,t)C(t,t')(I+G^T)^{-1}(t',1) \\ = G(1,0)^2 - 2C(1,0)G(1,0) + C(1,0) \quad (63)$$

From the above formulae we see that if the initial condition is large, in particular for $y(0) \gg \sqrt{\alpha}$, we have

$$\lim_{n_s \rightarrow \infty} C(1,0) = 1 \quad \lim_{n_s \rightarrow \infty} G(1,0) = 0 \quad (64)$$

$$\lim_{n_s \rightarrow \infty} L(1,0) = 1 \quad \lim_{n_s \rightarrow \infty} L(1,1) = 1 \quad (65)$$

for $0 < \lambda < 1$, and

$$\lim_{n_s \rightarrow \infty} C(1,0) = 0 \quad \lim_{n_s \rightarrow \infty} G(1,0) = 0 \quad (66)$$

$$\lim_{n_s \rightarrow \infty} L(1,0) = 0 \quad \lim_{n_s \rightarrow \infty} L(1,1) = 1 \quad (67)$$

for $\lambda > 1$, provided n_p is finite. The latter limits indicate that, as is to be expected, for $\lambda > 1$ the agent de-activates immediately after the first time step and starts being active and inactive alternatively. The former limits imply instead that the effective agent continues to play. In particular, as long as he's playing, he's insensitive to small perturbations, so that

$$y(t) \simeq (1-\lambda)y(0) + t\sqrt{\alpha}z(0) \quad (68)$$

Hence one sees that the de-freezing occurs for times of the order of $(1-\lambda)/\sqrt{\alpha}$.

In order to obtain a deeper insight on the stationary states of (33), we now turn to a different approach, namely the Eissfeller-Opper method. In a nutshell, the idea is to simulate many copies of the effective dynamics and calculate relevant physical observables as averages over the whole population. The core of the procedure lies in the possibility of evaluating the response function without actually adding an external probing field. In fact, for $G(t,t')$, which is formally given by $\langle \frac{\partial n(t)}{\partial h(t')} \rangle$, one can again resort to the noise and write

$$G(t,t') = \frac{1}{\sqrt{\alpha}} \langle \frac{\partial n(t)}{\partial z(t')} \rangle \\ = \frac{1}{\sqrt{\alpha}} \int \frac{\partial n(t)}{\partial z(t')} P(z) Dz \quad (69)$$

Now the noise distribution $P(z)$ is

$$P(z) \sim \exp \left[-\frac{1}{2} \sum_{t,t'} z(t)L^{-1}(t,t')z(t') \right] \quad (70)$$




FIG. 11: GCMG with different values of λ : numerical solution of the effective dynamics for $M = 250000$ copies with $\epsilon = -0.1$. Markers correspond to the volatility as a function of time for $n_s = 0.1$ (circles), 1 (squares), 2 (diamonds), 4 (up triangles), 8 (down triangles). Only a few markers are shown for simplicity.

so that, after an integration by parts, one gets

$$G(t,t') = \sum_{t''=0}^{t-1} \langle n(t)z(t'') \rangle L^{-1}(t'',t') \\ \equiv \sum_{t''=0}^{t-1} K(t,t'')L^{-1}(t'',t') \quad (71)$$

The matrix K , as well as the correlation function C , can be evaluated by an average over the copies (say, M) of the effective dynamics:

$$K(t,t') = \frac{1}{M} \sum_{\ell=1}^M n_{\ell}(t)z_{\ell}(t') \quad (72)$$

$$C(t,t') = \frac{1}{M} \sum_{\ell=1}^M n_{\ell}(t)n_{\ell}(t') \quad (73)$$

and the only remaining problem is that of generating a noise $z(t)$ having the desired statistical properties. This can be done by properly summing and re-scaling unit Gaussian variables. We focused again on the time-evolution of the quantity $L(t,t)$, whose limit $t \rightarrow \infty$ is linked to the volatility (see (55)). Results for different λ 's and n_s at $n_p = 1$ are shown in Fig. (11). One sees that finite memory generally increases fluctuations. The effect is more pronounced for large values of n_s .

VI. CONCLUSIONS

To summarize, we have studied the effects induced by a finite score memory in MG-based market models. Our main result is that finite score memory does not destroy market-like phenomenology in grand canonical minority

games, and remedies two embarrassing problems of minority games with a fixed set of strategies. From this point of view, implementing finite scores is a useful extension of GCMG, but cannot be understood fully by current analytical methods.

References

-
- [1] D. Challet and Y.-C. Zhang, *Physica A* **246** 407 (1997)
 - [2] D. Challet, M. Marsili and Y.-C. Zhang, *Minority Games and beyond*, Oxford University Press (Oxford, UK, 2004), forthcoming.
 - [3] D. Challet, M. Marsili and R. Zecchina, *Phys. Rev. Lett.* **84** 1824 (2000)
 - [4] A. De Martino and M. Marsili, *J. Phys. A: Math Gen.* **34** 2525 (2001)
 - [5] J.A.F. Heimerl and A.C.C. Coolen, *Phys. Rev. E* **63** 056121 (2001)
 - [6] J.A.F. Heimerl and A. De Martino, *J. Phys. A: Math Gen.* **34** L539 (2001)
 - [7] A.C.C. Coolen and J.A.F. Heimerl, *J. Phys. A: Math Gen.* **34** 10783 (2001)
 - [8] M. Marsili, R. Mulet, F. Ricci-Tersenghi and R. Zecchina, *Phys. Rev. Lett.* **87** 208701 (2001)
 - [9] D. Challet, A. De Martino and M. Marsili, *Physica A* **338** 143 (2004)
 - [10] P. Jefferies, M.L. Hart, P.M. Hui and N.F. Johnson, *Eur. Phys. J. B* **20** 493 (2001)
 - [11] M.L. Hart, D. Lamper and N.F. Johnson, *Physica A* **316** 649 (2002)
 - [12] D. Challet and M. Marsili, *Phys. Rev. E* **68** 036132 (2003)
 - [13] T. Galla, unpublished.
 - [14] H. Eissfeller and M. Opper, *Phys. Rev. Lett.* **68** 2094 (1992)
 - [15] D. Challet, M. Marsili and Y.-C. Zhang, *Physica A* **299** 228 (2001)
 - [16] D. Challet, M. Marsili and Y.-C. Zhang, *Physica A* **276**, 284 (2000)
 - [17] J. P. Garrahan, E. Moro, D. Sherrington, *Phys. Rev. E* **62** R9 (2000)
 - [18] J. D. Farmer, *Industrial and Corporate Change* **11** 895 (2002); SFI Working Paper 98-12-117
 - [19] P. Jefferies, M.L. Hart, P.M. Hui and N.F. Johnson, *Int. J. Th and Appl. Fin.* **3** 3 (2000).
 - [20] J.-Ph. Bouchaud and M. Potters, *Theory of financial risks*, Cambridge University Press (Cambridge, UK, 2000)
 - [21] M. M. Dacorogna, R. Gençay, U. Müller, R. B. Olsen and O. V. Pictet, *An introduction to high-frequency finance*, Academic Press (London, UK, 2001)
 - [22] M. Marsili, D. Challet, *Phys. Rev. E* **64** 056138 (2001)
 - [23] D. Challet and M. Marsili, *Phys. Rev. E* **60** R6271 (1999)

APPENDIX A: ORIGINAL MG WITH $\lambda > 0$.

The definition of the original MG is similar to that of the GCMG, except that all the agents have at least two strategies, and one score per strategy. Assuming that they have only two strategies, denoted $a_{i,1}$ and $a_{i,2}$ for agent i , each having their respective score $y_{i,1}$ and $y_{i,2}$.




FIG. 12: The same realization of the standard minority game with unbiased initial strategy scores and $\lambda = 0$ (circles), biased initial scores and $\lambda = 0$ (squares), biased initial scores and $\lambda = 0.1$ (triangles). ($P = 100$, $\alpha = 0.1$)

The latter evolve according to

$$y_{i,s}(t+1) = y_i(t)(1 - \lambda/P) - a_{i,s}^{\mu(t)} A(t) \quad (\text{A1})$$

for the on-line MG, and

$$y_{i,s}(t+1) = y_i(t)(1 - \lambda) - a_{i,s}^{\mu(t)} A(t) \quad (\text{A2})$$

for the batch MG. At each time step, the agents use their strategy that has the highest score. The original MG was defined with $\lambda = 0$, and the usual control parameter is $\alpha = P/N$. There is a phase transition at $\alpha_c = 0.3374\dots$ that separates an asymmetric, information rich phase with $H > 0$ ($\alpha > \alpha_c$) from a symmetric phase with $H = 0$ ($\alpha < \alpha_c$). In the latter, the stationary state depends on the asymmetry

$$U_0 = U_{i,1}(0) - U_{i,2}(0)$$

of the initial conditions. This dependence is due to the fact that, if $\lambda = 0$, all market fluctuations since time $t = 0$ are remembered and contribute with the same weight to $U_{i,s}(t)$, irrespective of how far in the past they took place.

Fig. 12 shows that a time dependent $\sigma^2 = \langle A^2 \rangle_t$ (here, $\langle \dots \rangle_t$ stands for an average over a long but finite time interval around t). converges, for long times, to a value




FIG. 13: Average volatility $\sigma^2/N = \langle A^2 \rangle / N$ versus $\alpha = P/N$ of the standard MG with unbiased initial scores and $\lambda = 0$ (circles), $\lambda = 0.1$ (squares), $\lambda = 0.2$ (diamonds), with biased initial scores and $\lambda = 0.2$ (triangles).




FIG. 14: σ^2/N (black lines) and H/N (red lines) for a realization of the standard Minority Game with increasing λ

which is independent of initial conditions U_0 . This implies that the initial asymmetry U_0 has no influence any more on the stationary state.

This is confirmed by Fig. 13, where we plot σ^2/N as a function of α for various U_0 and λ . For a fixed α and for a given realization of the game, σ^2 and H are increasing functions of λ , although in the symmetric phase ($\alpha < \alpha_c = 0.3374\dots$) σ^2 varies very slowly with λ (see Fig. 14). While only small values of λ make physical sense, for completeness the same figure reports also the strange effects of very large λ . From (A1), $\lambda = P$ cancels the contribution of $U_i(t)$ in $U_i(t+1)$; larger λ implies that this contribution is of opposite sign to the usual MG.

List of other working papers:

2004

1. Xiaohong Chen, Yanqin Fan and Andrew Patton, Simple Tests for Models of Dependence Between Multiple Financial Time Series, with Applications to U.S. Equity Returns and Exchange Rates, WP04-19
2. Valentina Corradi and Walter Distaso, Testing for One-Factor Models versus Stochastic Volatility Models, WP04-18
3. Valentina Corradi and Walter Distaso, Estimating and Testing Sochastic Volatility Models using Realized Measures, WP04-17
4. Valentina Corradi and Norman Swanson, Predictive Density Accuracy Tests, WP04-16
5. Roel Oomen, Properties of Bias Corrected Realized Variance Under Alternative Sampling Schemes, WP04-15
6. Roel Oomen, Properties of Realized Variance for a Pure Jump Process: Calendar Time Sampling versus Business Time Sampling, WP04-14
7. Richard Clarida, Lucio Sarno, Mark Taylor and Giorgio Valente, The Role of Asymmetries and Regime Shifts in the Term Structure of Interest Rates, WP04-13
8. Lucio Sarno, Daniel Thornton and Giorgio Valente, Federal Funds Rate Prediction, WP04-12
9. Lucio Sarno and Giorgio Valente, Modeling and Forecasting Stock Returns: Exploiting the Futures Market, Regime Shifts and International Spillovers, WP04-11
10. Lucio Sarno and Giorgio Valente, Empirical Exchange Rate Models and Currency Risk: Some Evidence from Density Forecasts, WP04-10
11. Ilias Tsiakas, Periodic Stochastic Volatility and Fat Tails, WP04-09
12. Ilias Tsiakas, Is Seasonal Heteroscedasticity Real? An International Perspective, WP04-08
13. Damin Challet, Andrea De Martino, Matteo Marsili and Isaac Castillo, Minority games with finite score memory, WP04-07
14. Basel Awartani, Valentina Corradi and Walter Distaso, Testing and Modelling Market Microstructure Effects with an Application to the Dow Jones Industrial Average, WP04-06
15. Andrew Patton and Allan Timmermann, Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity, WP04-05
16. Andrew Patton, Modelling Asymmetric Exchange Rate Dependence, WP04-04
17. Alessio Sancetta, Decoupling and Convergence to Independence with Applications to Functional Limit Theorems, WP04-03
18. Alessio Sancetta, Copula Based Monte Carlo Integration in Financial Problems, WP04-02
19. Abhay Abhayankar, Lucio Sarno and Giorgio Valente, Exchange Rates and Fundamentals: Evidence on the Economic Value of Predictability, WP04-01

2002

1. Paolo Zaffaroni, Gaussian inference on Certain Long-Range Dependent Volatility Models, WP02-12
2. Paolo Zaffaroni, Aggregation and Memory of Models of Changing Volatility, WP02-11
3. Jerry Coakley, Ana-Maria Fuertes and Andrew Wood, Reinterpreting the Real Exchange Rate - Yield Diffential Nexus, WP02-10
4. Gordon Gemmill and Dylan Thomas , Noise Training, Costly Arbitrage and Asset Prices: evidence from closed-end funds, WP02-09
5. Gordon Gemmill, Testing Merton's Model for Credit Spreads on Zero-Coupon Bonds, WP02-08
6. George Christodoulakis and Steve Satchell, On th Evolution of Global Style Factors in the MSCI Universe of Assets, WP02-07
7. George Christodoulakis, Sharp Style Analysis in the MSCI Sector Portfolios: A Monte Caro Integration Approach, WP02-06

8. George Christodoulakis, Generating Composite Volatility Forecasts with Random Factor Betas, WP02-05
9. Claudia Riveiro and Nick Webber, Valuing Path Dependent Options in the Variance-Gamma Model by Monte Carlo with a Gamma Bridge, WP02-04
10. Christian Pedersen and Soosung Hwang, On Empirical Risk Measurement with Asymmetric Returns Data, WP02-03
11. Roy Batchelor and Ismail Orgakcioglu, Event-related GARCH: the impact of stock dividends in Turkey, WP02-02
12. George Albanis and Roy Batchelor, Combining Heterogeneous Classifiers for Stock Selection, WP02-01

2001

1. Soosung Hwang and Steve Satchell , GARCH Model with Cross-sectional Volatility; GARCHX Models, WP01-16
2. Soosung Hwang and Steve Satchell, Tracking Error: Ex-Ante versus Ex-Post Measures, WP01-15
3. Soosung Hwang and Steve Satchell, The Asset Allocation Decision in a Loss Aversion World, WP01-14
4. Soosung Hwang and Mark Salmon, An Analysis of Performance Measures Using Copulae, WP01-13
5. Soosung Hwang and Mark Salmon, A New Measure of Herding and Empirical Evidence, WP01-12
6. Richard Lewin and Steve Satchell, The Derivation of New Model of Equity Duration, WP01-11
7. Massimiliano Marcellino and Mark Salmon, Robust Decision Theory and the Lucas Critique, WP01-10
8. Jerry Coakley, Ana-Maria Fuertes and Maria-Teresa Perez, Numerical Issues in Threshold Autoregressive Modelling of Time Series, WP01-09
9. Jerry Coakley, Ana-Maria Fuertes and Ron Smith, Small Sample Properties of Panel Time-series Estimators with I(1) Errors, WP01-08
10. Jerry Coakley and Ana-Maria Fuertes, The Felsdtein-Horioka Puzzle is Not as Bad as You Think, WP01-07
11. Jerry Coakley and Ana-Maria Fuertes, Rethinking the Forward Premium Puzzle in a Non-linear Framework, WP01-06
12. George Christodoulakis, Co-Volatility and Correlation Clustering: A Multivariate Correlated ARCH Framework, WP01-05
13. Frank Critchley, Paul Marriott and Mark Salmon, On Preferred Point Geometry in Statistics, WP01-04
14. Eric Bouyé and Nicolas Gaussel and Mark Salmon, Investigating Dynamic Dependence Using Copulae, WP01-03
15. Eric Bouyé, Multivariate Extremes at Work for Portfolio Risk Measurement, WP01-02
16. Erick Bouyé, Vado Durrleman, Ashkan Nikeghbali, Gael Riboulet and Thierry Roncalli, Copulas: an Open Field for Risk Management, WP01-01

2000

1. Soosung Hwang and Steve Satchell , Valuing Information Using Utility Functions, WP00-06
2. Soosung Hwang, Properties of Cross-sectional Volatility, WP00-05
3. Soosung Hwang and Steve Satchell, Calculating the Miss-specification in Beta from Using a Proxy for the Market Portfolio, WP00-04
4. Laun Middleton and Stephen Satchell, Deriving the APT when the Number of Factors is Unknown, WP00-03
5. George A. Christodoulakis and Steve Satchell, Evolving Systems of Financial Returns: Auto-Regressive Conditional Beta, WP00-02
6. Christian S. Pedersen and Stephen Satchell, Evaluating the Performance of Nearest Neighbour Algorithms when Forecasting US Industry Returns, WP00-01

1999

1. Yin-Wong Cheung, Menzie Chinn and Ian Marsh, How do UK-Based Foreign Exchange Dealers Think Their Market Operates?, WP99-21
2. Soosung Hwang, John Knight and Stephen Satchell, Forecasting Volatility using LINEX Loss Functions, WP99-20
3. Soosung Hwang and Steve Satchell, Improved Testing for the Efficiency of Asset Pricing Theories in Linear Factor Models, WP99-19
4. Soosung Hwang and Stephen Satchell, The Disappearance of Style in the US Equity Market, WP99-18
5. Soosung Hwang and Stephen Satchell, Modelling Emerging Market Risk Premia Using Higher Moments, WP99-17
6. Soosung Hwang and Stephen Satchell, Market Risk and the Concept of Fundamental Volatility: Measuring Volatility Across Asset and Derivative Markets and Testing for the Impact of Derivatives Markets on Financial Markets, WP99-16
7. Soosung Hwang, The Effects of Systematic Sampling and Temporal Aggregation on Discrete Time Long Memory Processes and their Finite Sample Properties, WP99-15
8. Ronald MacDonald and Ian Marsh, Currency Spillovers and Tri-Polarity: a Simultaneous Model of the US Dollar, German Mark and Japanese Yen, WP99-14
9. Robert Hillman, Forecasting Inflation with a Non-linear Output Gap Model, WP99-13
10. Robert Hillman and Mark Salmon, From Market Micro-structure to Macro Fundamentals: is there Predictability in the Dollar-Deutsche Mark Exchange Rate?, WP99-12
11. Renzo Avesani, Giampiero Gallo and Mark Salmon, On the Evolution of Credibility and Flexible Exchange Rate Target Zones, WP99-11
12. Paul Marriott and Mark Salmon, An Introduction to Differential Geometry in Econometrics, WP99-10
13. Mark Dixon, Anthony Ledford and Paul Marriott, Finite Sample Inference for Extreme Value Distributions, WP99-09
14. Ian Marsh and David Power, A Panel-Based Investigation into the Relationship Between Stock Prices and Dividends, WP99-08
15. Ian Marsh, An Analysis of the Performance of European Foreign Exchange Forecasters, WP99-07
16. Frank Critchley, Paul Marriott and Mark Salmon, An Elementary Account of Amari's Expected Geometry, WP99-06
17. Demos Tambakis and Anne-Sophie Van Royen, Bootstrap Predictability of Daily Exchange Rates in ARMA Models, WP99-05
18. Christopher Neely and Paul Weller, Technical Analysis and Central Bank Intervention, WP99-04
19. Christopher Neely and Paul Weller, Predictability in International Asset Returns: A Re-examination, WP99-03
20. Christopher Neely and Paul Weller, Intraday Technical Trading in the Foreign Exchange Market, WP99-02
21. Anthony Hall, Soosung Hwang and Stephen Satchell, Using Bayesian Variable Selection Methods to Choose Style Factors in Global Stock Return Models, WP99-01

1998

1. Soosung Hwang and Stephen Satchell, Implied Volatility Forecasting: A Comparison of Different Procedures Including Fractionally Integrated Models with Applications to UK Equity Options, WP98-05
2. Roy Batchelor and David Peel, Rationality Testing under Asymmetric Loss, WP98-04
3. Roy Batchelor, Forecasting T-Bill Yields: Accuracy versus Profitability, WP98-03
4. Adam Kurpiel and Thierry Roncalli, Option Hedging with Stochastic Volatility, WP98-02
5. Adam Kurpiel and Thierry Roncalli, Hopscotch Methods for Two State Financial Models, WP98-01