



Modeling market mechanism with minority game

Damien Challet^{a,*}, Matteo Marsili^b, Yi-Cheng Zhang^a

^a*Institut de Physique Théorique, Université de Fribourg, 1700 Fribourg, Switzerland*

^b*Istituto Nazionale per la Fisica della Materia (INFN), Trieste-SISSA Unit, V. Beirut 2-4, Trieste I-34014, Italy*

Received 17 September 1999

Abstract

Using the minority game model we study a broad spectrum of problems of market mechanism. We study the role of different types of agents: producers, speculators as well as noise traders. The central issue here is the information flow: producers feed in the information whereas speculators make it away. How well each agent fares in the common game depends on the market conditions, as well as their sophistication. Sometimes there is much to gain with little effort, sometimes great effort virtually brings no more incremental gain. Market impact is also shown to play an important role, a strategy should be judged when it is actually used in play for its quality. Though the minority game is an extremely simplified market model, it allows to ask, analyze and answer many questions which arise in real markets. © 2000 Elsevier Science B.V. All rights reserved.

1. Introduction

Recently, it became possible to study markets of heterogeneous agents, in particular in the form of the so-called minority games (MG) [1,2]. Since long-time practitioners of the market, as well as some economists have criticized the main-stream economics where a so-called representative agent plays the central role. Many prominent economists like Herbert Simon [3], Richard Day and Brian Arthur [4] have been forceful proponents of the “bounded rationality” and “inductive thinking”. However, though many people join in unison in their criticism of the main stream, their alternative approaches and models do not command consensus yet.

The MG is inspired by Arthur’s “*El Farol*” model [4], which shows for the first time how the equilibrium can be reached using inductive thinking. Whereas *El Farol* model is about the equilibrium, our MG model is about fluctuations. In a sense MG

* Corresponding author. Fax: +41-26-300-9758.

E-mail address: damien.challet@unifr.ch (D. Challet)

gives us a powerful tool to study detailed pattern of fluctuations, the equilibrium point is trivial by design. It is the fluctuations that play the dominant role in economic activities, like the market mechanism. MG allows us to study in a precise manner the approach to equilibrium, how the agents try to outsmart each other, for their selfish gain, compete for the available marginal information (any deviation from the mid-point represents exploitable advantage). It is for this residual margin that all the agents fight for, resembling the real markets. The importance to market mechanism is primordial, as any practitioner can attest. Neo-classical economics would tell us that the MG, as in a competitive market, does not offer consistent gain, based on the efficient market hypothesis (EMH). However, if some agents stop playing (i.e. choosing dynamically among the two sides), they will give away information that the other more diligent dynamic agents make use of. This means that equilibrium can be only dynamically maintained, any relaxing would imply a relative disadvantage. It is not that the same thing occurs in real markets?

Studying a model of market mechanism opens up many detailed questions, which practitioners have to face constantly but the main-stream economists do not have any clue to answer them. For instance, in a model like MG agents interact with each other through a common market, what information each agent brings in? What gain each agent takes out? How sophisticated should an agent be? What is realizable gain objective? What is the role of noise traders? How about insider trading (an agent processing privileged information about fellow agents)? What is the market impact of an otherwise clever strategy? The list is obviously endless. The point we want to make here is that with so little to start with, and with so many questions relevant to real markets one can hope for a qualitative answer.

After two years since the MG's birth, during which much work has revealed its extremely rich structure [5], an analytical approach leading to its exact solution has been found [6,7]. Unfortunately, the main progress is still confined in the physics community. We hope that, with this paper, this will change: The aim is to convince people, including economists hopefully, that many concrete questions about market mechanisms can be asked and answered, in precise and analytical way, using the approach of Refs. [6,7]. In fact the MG can be used as a flexible platform and different handles can be added and manipulated almost at will. To achieve our goal, the analytic approach shall be supplemented by numerical simulations to confirm its validity. More technical parts and heavy calculations shall be dealt with in the appendices.

2. Main results

Here below we give a list of salient points of our paper:

(1) *Diversification of ideas.* If an agent has different alternative strategies, it is better to have them diversified, i.e., not too much correlated. We show the effects of diversification.

(2) *Markets have two types of agents: producers and speculators.* The former do not have alternative strategies; the latter are represented by the normal agents of the standard MG. Producers provide information into the market, upon which speculators feed. For the first time it is possible to demonstrate that producers and speculators need each other, they live in a symbiosis. However, benefits to each group are not equal, depending on the parameters.

(3) *Agents are not obliged to play, if they do not see a possible gain.* We generalize MG to let agents have the option of not playing. In the presence of producers, markets appear to be attractive and more speculators are drawn into the fore.

(4) *Noise traders.* One may wonder if some traders decide to be pure noise traders, i.e., they use completely random strategy, what is the “harm” done to other market participants (producers and speculators), as well as to themselves. In the information rich phase, they appear to increase volatility and in the herding-effect phase, they actually make the market perform better.

(5) Despite the fact that agents start equally equipped, there are better and worse agents and the rank of the agents has an interesting non-Gaussian “bar-code” structure.

(6) Sometimes it pays to increase the capacity of an agent’s brain, say add one more unit in M . This will give enormous advantage to the better-equipped agent in the crowded phase (or symmetric phase), where information on the range M is exhausted, whereas such a feature becomes a disadvantage in the information-rich phase.

(7) *Does it pay to have more strategies as alternatives? In general yes.* Here we calculate the relative advantage by having more alternatives. We also show that, due to self-market impact, the imagined gain differs from the real gain, a fact known too well to market practitioners. Even each agent has many alternatives, they actually use only a small number of them.

(8) *Some agents may get illegal information about others.* It is just like a stock broker who knows his clients’ orders before execution. Hence, he has privileged information and should be barred from trading. An agent who spies on fellow agents enjoys trading advantages. We measure how much is this effect, as the number of fellow agents whom you spy increases, how much would be your gain.

3. Formalism and review

Our model of market consists of N agents which, for simplicity, can take only one of two actions, such as “buy” and “sell” at each time step t . We represent this assuming that each agent $i=1, \dots, N$, at time t , can either do the action $a_i(t)=+1$ or the opposite action $a_i(t)=-1$. Given the actions of all agents, the gain of agent i is given by

$$g_i(t) = -a_i(t)A(t) \quad \text{where} \quad A(t) = \sum_{j=1}^N a_j(t). \quad (1)$$

This equation models the basic structure of market interaction where each agent’s payoffs are determined by the action taken and by a global quantity $A(t)$, which is

usually a price and it is determined by all of them. For simplicity, we assume here a linear dependence of $g_i(t)$ on $A(t)$. Other choices, such as $g_i(t) = -a_i(t) \text{sign} A(t)$ in Refs. [1,8,9], can be taken without affecting qualitatively the results we shall discuss below. This interaction clearly rewards the minority of agents (those who took the action $a_i(t) = -\text{sign} A(t)$) who gain an amount $|A(t)|$ and punishes the majority by a loss $-|A(t)|$, hence the name minority game [1]. There are always more losers than winners and agents have no way of knowing what the majority will do before taking their actions.

All agents have access to public information which is represented by an integer variable μ taking one of P values. At time t information “takes the value” $\mu(t)$. We shall also call $\mu(t)$ history since originally this information has been introduced as encoding the record of the past $M = \log_2 P$ signs of $A(t)$ with M bits. It has however been shown [10] that if $\mu(t)$ is randomly drawn in $\{1, \dots, P\}$ one recovers the same results (see also the discussion in Refs. [6,7]). We shall henceforth consider this second, simpler case. When having access to some information, agents can behave differently for different values of $\mu(t)$, eventually because of their personal beliefs on the impact that information $\mu(t)$ shall have on the outcome of the market, $A(t)$. Strictly speaking $A(t)$ only depends on what agents do, so $\mu(t)$ has no direct impact on the market. However, if agents behavior depends on $\mu(t)$ also $A(t)$ shall depend on it, and we denote it by $A^{\mu(t)}(t)$.

How do agents choose actions under information $\mu(t)$? If agents expect that $\mu(t)$ contains some information on the market, they will consider *forecasting strategies* which for each value of μ suggest which action a^μ shall be done. There are 2^P such strategies, and we assume, for the time being, that each agent just picks S such rules randomly (with replacement) from the set of all 2^P strategies. The action of agent i if she follows her s th strategy and the information is μ is denoted by $a_{s,i}^\mu$. Therefore, if $s_i(t)$ is the choice made (in a way we shall specify below) by agent i at time t , her action becomes $a_i(t) \rightarrow a_{s_i(t),i}^{\mu(t)}$ and correspondingly, her gain [Eq. (1)] becomes

$$g_i(t) = -a_{s_i(t),i}^{\mu(t)} A^{\mu(t)}(t) \quad \text{where} \quad A^{\mu(t)}(t) = \sum_{j=1}^N a_{s_j(t),j}^{\mu(t)}. \tag{2}$$

In this paper, we mainly focus on $S = 2$. This case contains all the richness of the model and allows a more transparent presentation. All the results discussed below can be extended to $S > 2$ along the lines of Ref. [7]. For $S = 2$ we can adopt a notation where each agent controls a variable $s_i \in \{\downarrow, \uparrow\}$, with the identification $\uparrow = +1$ and $\downarrow = -1$. This is useful to distinguish strategies s_i from actions a_i . It is convenient to introduce the variables

$$\omega_i^\mu = \frac{a_{\uparrow,i}^\mu + a_{\downarrow,i}^\mu}{2}, \quad \zeta_i^\mu = \frac{a_{\uparrow,i}^\mu - a_{\downarrow,i}^\mu}{2}. \tag{3}$$

With these notations, the action taken by this agent in reaction to the history μ is

$$a_{i,s_i}^\mu = \omega_i^\mu + \zeta_i^\mu s_i, \tag{4}$$

so ω_i^μ represents the part of i 's strategies which is fixed, whereas ξ_i^μ is the variable part. We also define $\Omega^\mu = \sum_i \omega_i^\mu$ so that

$$A^\mu(t) = \sum_{i=1}^N a_{i,s_i(t)}^\mu = \Omega^\mu + \sum_{i=1}^N \xi_i^\mu s_i(t). \tag{5}$$

Each agent updates the cumulated virtual payoffs of all her strategies according to

$$U_{s,i}(t + 1) = U_{s,i}(t) - A^{\mu(t)}(t) a_{i,s}^{\mu(t)} \tag{6}$$

The quantity $U_{i,s}$ is a “reliability index” which quantifies the agent i 's perception of the success of her s th strategy. $U_{i,s}(t)$ is the *virtual* cumulated payoff that agent i would have received up to time t if she had always played strategy s (with others playing the strategies $s_j(t')$ which they actually played at times $t' < t$). *Virtual* here means that this is not the real cumulated payoff but rather that *perceived* by agent i . These differ, as explained below and in Ref. [7], because agents neglect their impact on the market (i.e., the fact that if they had indeed always played s the aggregate quantity $A(t)$ would have been different).

Inductive dynamics [4,1] consists in assuming that agents trust and use their most reliable strategy, which are those with the largest virtual score:

$$s_i(t) = \arg \max_{s \in \{\uparrow, \downarrow\}} U_{i,s}(t). \tag{7}$$

More generally, one can consider a probabilistic choice rule – the so called *Logit* model [11] – such that $P(s_i(t) = s) \propto \exp[\Gamma U_{s,i}(t)]$, (see [12,6,7]). Then Eq. (7) is recovered in the limit $\Gamma \rightarrow \infty$. As in Ref. [13], we find it useful to introduce the variables $\Delta_i(t) = U_{i,\uparrow} - U_{i,\downarrow}$. Their dynamics reads

$$\Delta_i(t + 1) = \Delta_i(t) - A^{\mu(t)}(t) \xi_i^{\mu(t)} \tag{8}$$

and Eq. (7) becomes

$$s_i(t) = \text{sign } \Delta_i(t). \tag{9}$$

3.1. Notations on averages

We define the temporal average of a given time-dependent quantity $R(t)$ as

$$\langle R \rangle = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T R(t). \tag{10}$$

This quantity can be decomposed into conditional averages on histories, that is

$$\langle R^\mu \rangle = \lim_{T \rightarrow \infty} \frac{P}{T} \sum_{t=1}^T R(t) \delta_{\mu(t),\mu}. \tag{11}$$

Note that the factor P and the relation $\langle \delta_{\mu(t),\mu} \rangle = 1/P$ imply that $\langle R^\mu \rangle$ is a conditional average. More precisely, it is the temporal average of the quantity $R(t)$ subject to the

condition that the actual history $\mu(t)$ was¹ μ . Finally, averages over the histories μ of a quantity R^μ are defined as

$$\bar{R} \equiv \frac{1}{P} \sum_{\mu=1}^P R^\mu . \tag{12}$$

3.2. Quantities of interest

With these notations, let us now discuss the main quantities which characterize the stationary state of the system. The main free parameter, as first observed in Ref. [8], is

$$\alpha = \frac{P}{N} \tag{13}$$

and we shall eventually consider the thermodynamic limit where $N, P \rightarrow \infty$ with α fixed. The first quantity of interest is

$$\sigma^2 \equiv \overline{\langle A \rangle^2} = \bar{\Omega}^2 + 2 \sum_{i=1}^N \overline{\Omega \xi_i} \langle s_i \rangle \langle s_i \rangle + \sum_{i,j} \overline{\xi_i \xi_j} \langle s_i s_j \rangle . \tag{14}$$

This equals the total loss of agents

$$- \sum_i \langle g_i \rangle = \sigma^2 , \tag{15}$$

so it is a measure of global waste. It also quantifies the volatility of the market, i.e., the fluctuations of the quantity $A(t)$, and is related to the average “distance” between agents (see Appendix A). Even though $\langle A \rangle = 0$, by symmetry, it may happen that for a particular μ , the aggregate quantity $A(t)$ is nonzero on average, i.e., that $\langle A^\mu \rangle \neq 0$. In order to quantify this asymmetry, we introduce the quantity

$$H \equiv \overline{\langle A \rangle^2} = \bar{\Omega}^2 + 2 \sum_{i=1}^N \overline{\Omega \xi_i} \langle s_i \rangle \langle s_i \rangle + \sum_{i,j} \overline{\xi_i \xi_j} \langle s_i \rangle \langle s_j \rangle . \tag{16}$$

Note that the only difference with σ^2 lies in the diagonal terms ($i = j$) of the last sum. Indeed, we assume that $\langle s_i s_j \rangle = \langle s_i \rangle \langle s_j \rangle$ for $i \neq j$, whereas² $\langle s_i^2 \rangle \equiv 1 \neq \langle s_i \rangle^2$. Indeed, we can write

$$\sigma^2 = H + \sum_{i=1}^N \overline{\xi_i^2} (1 - \langle s_i \rangle^2) . \tag{17}$$

If $H > 0$, the game is asymmetric: At least for some μ one has that $\langle A^\mu \rangle \neq 0$. This implies that there is a *best* strategy $a_{\text{best}}^\mu = -\text{sign} \langle A^\mu \rangle$ which in principle could give a

¹ This implies that the number of iterations must be proportional to P in any numerical simulation.
² This amounts to say that the fluctuations in time of s_i around its average $\langle s_i \rangle$ are uncorrelated with $s_j - \langle s_j \rangle$. This assumption fails when crowd effects occur, i.e., in the symmetric phase, and our theory will accordingly fail to describe these effects.

positive gain $|\overline{\langle A \rangle}| - 1$.³ In economic terms we may say that the system is not *arbitrage free*, and that H is a measure of the perceived arbitrage opportunities present in the market. As a function of $\alpha = P/N$ the system displays a *phase transition* with symmetry breaking [13]: For $\alpha > \alpha_c$ the symmetry between the two signs of $A(t)$ is broken.

H plays a particular important role because in Refs. [6,7] it has been shown that the inductive dynamics is equivalent to a dynamics which minimizes H in the dynamical variables $m_i = \langle s_i \rangle$. Therefore, the ground-state properties of the Hamiltonian H yields the stationary state of the system. H is a spin-glass Hamiltonian where $\overline{\Omega \xi_i}$ are the local magnetic fields and $\overline{\xi_i \xi_j}$ the coupling between two agents. These play the same role as *quenched* disorder in spin glasses. This system is of *mean field* type since interactions $\overline{\xi_i \xi_j}$ are infinite ranged. For this reason, the statistical mechanics approach to disordered systems [14,15] via the replica method yields exact results for these models (see Appendix C).

The behavior of each agent is completely determined by the difference of her cumulated payoffs Δ_i . For long times, $\Delta_i \simeq v_i t$, where

$$v_i = \langle \Delta(t+1) - \Delta_i(t) \rangle = -2 \overline{\langle A \rangle \xi_i} . \tag{18}$$

If $v_i \neq 0$, agent i shall stick to only one strategy $s_i = \text{sign } v_i$, whereas if $v_i = 0$, she will sometimes use her \uparrow strategy and sometimes her \downarrow one. This is quantified by m_i , and a global measure of the fluctuations in the strategic choices of agents is given by

$$Q = \frac{1}{N} \sum_{i=1}^N m_i^2 . \tag{19}$$

This quantity also emerges naturally from the replica approach where it plays a key role.

4. Speculators with diversified strategies

In the standard MG, it is assumed that the agents draw all their strategies randomly, and independently. One can argue that the agents can be less simple-minded so that they first draw a strategy, and then following their needs or what seems the best for them, draw the others strategies. For instance, if $S = 2$, an agent can believe that one strategy is enough and sticks to it (or takes two same strategies). On the contrary, an agent might believe that it is better to have one strategy and another one which is quite opposite. More generally, we suppose that all the agents⁴ draw their second strategy according to

$$P(a_{\uparrow}^{\mu} = a_{\downarrow}^{\mu}) = c \quad \forall \mu . \tag{20}$$

³ Here the -1 comes from the fact that if the strategy is actually played $A^{\mu} \rightarrow A^{\mu} + a_{\text{best}}^{\mu}$ and “in principle” means that A^{μ} would also change as a result of the fact that other agents would also react to the best strategy agent.

⁴ This can be generalized to a c for each agent; exact results also arise from the replica calculus.

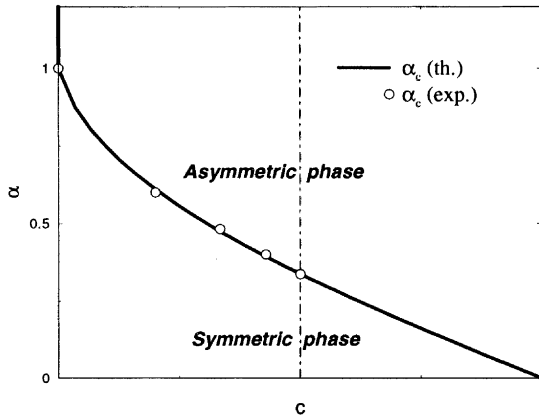


Fig. 1. Phase diagram of the minority game with diversified strategies. The phase transition in the standard MG corresponds to the dash-dotted vertical line $c = \frac{1}{2}$. The circle are numerical data.

The parameter c counts the average fraction of histories for which the agents' choices are biased, that is, the average correlation between their two strategies. The standard MG corresponds to the independent case $c = \frac{1}{2}$, while having only one strategy is obtained with $c = 1$. The other very special case is $c = 0$: all agents have two opposite strategies, thus there is no asymmetry in the outcome. As a result, the game is always in the symmetric phase: as α is varied, no phase transition occurs. Increasing c has two effects: on the one hand, it increases the bias of the outcome $\Omega^\mu \sim \sqrt{cN}$, on the other, it reduces the ability of the agents of being adaptative, since they learn something about the game only when $\zeta_i^\mu \neq 0$ (see Eq. (8)), which happens in average for $(1 - c)P$ histories. The fact that the biases depend on c too implies that the second-order phase transition also occurs when this parameter is varied. With the replica formalism (see Appendix C), one gets the phase diagram of the MG with parameter c (see Fig. 1). In the standard MG, one varies α (dot-dashed vertical line). If one fixes α and changes c , the symmetry is also broken (any horizontal line). Note that if $c = 0$ and $\alpha > 1$, an infinitesimal c breaks the symmetry of the game.

5. Speculators and producers

Real markets are not *zero sum* games [16]. The fact that most participants are interested in playing is beyond doubt. In real markets the participants can be grossly divided into two groups: speculators and producers [16]. Producers can be characterized by those using the market for purposes other than speculation. They need market for hedging, financing, or any ordinary business. They thus pay less or no attention to "timing the market". Speculators, on the other hand, join the market with the aim of exploiting the marginal profit pockets. The two groups were shown to live in symbiosis [16]: the former inject information into the market prices, and the latter make a living




Fig. 2. Gain of producers and speculators versus the number of producers (in P unit); the number of speculators is fixed at $N = 641$ ($c = 0$, $M = 8$, $S = 2$, $\alpha = 0.4$, average over 200 realizations). The lines are theoretical predictions.

carefully exploiting this information. One may wonder why do producers let themselves be taken advantage of. Our answer is that they have other, probably more profitable business in mind. To conduct their business, they need the market, and their expertises and talents in other areas give them still better games to play. Speculators, being less capable in other areas, or by choice, make do exploiting the “meager margin” left in the competitive market.

In our MG, these general questions can be studied in detail. Producers will be limited in choice, their activities outside the game are not represented. We define a *speculator* as a normal agent, and a *producer* as an agent limited to one strategy. Thus, the latter have a fixed pattern in their market behavior and put a measurable amount of information into the market, which is exploited by the speculators. We take a population of N speculators and always define $\alpha = P/N$. We add ρN heterogeneous producers, so that ρ is the fraction of producers per speculator. The outcome is then

$$A^\mu = A_{\text{spec}}^\mu + A_{\text{prod}}^\mu. \quad (21)$$

The bias induced by the producers adds to the one caused by the speculators, so that the total bias is of order $\sqrt{(c + \rho)N}$. Therefore, the phase transition can be obtained at fixed P by varying either N , c , or the number of producers. Let us begin with the last possibility. We fix $c = 0$, $P = 2^8$, $N = 641$ and plot the gains of the speculators and producers as a function of the number of producers (see Fig. 2). In the symmetric phase, the speculators wash out all the available information, thus, by symmetry, the gain of the producers (squares) is zero. As the number of producers increases, the gain of the speculators (circles) stays negative but grows monotonically, while the gain of the producers remains zero as long as the symmetry of the outcome is not broken. When the number of producers reaches a critical value, the speculators are

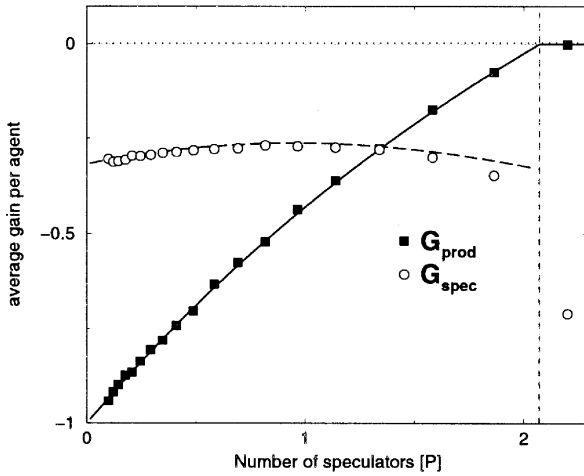


Fig. 3. Gain of producers and speculators versus the number of speculators (in P unit); the number of producers is fixed at 64 ($c = 0$, $M = 8$, $S = 2$, average over 200 realizations). The lines are theoretical predictions.

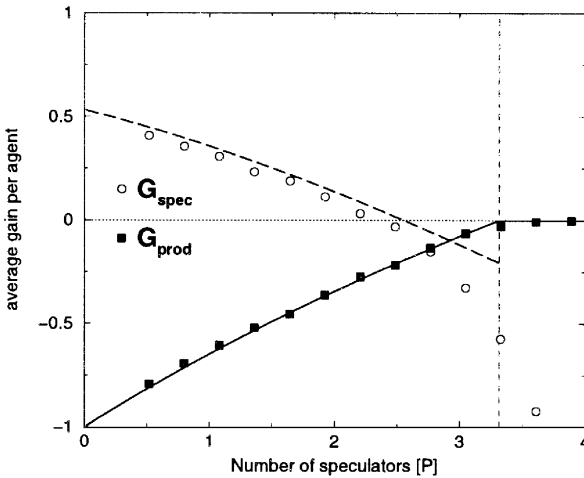


Fig. 4. Gain of producers and speculators versus the number of speculators (in P unit); the number of producers is fixed at 256 ($c = 0$, $M = 6$, $S = 2$, average over 200 realizations). The lines are theoretical predictions.

no more able to remove all the available information, therefore the (second-order) phase transition occurs (dashed line). Beyond this point, the producers lose more and more, while some (frozen) speculators gain more than zero in average (see Section 8). At one point, the gains of speculators and producers are the same. Finally, there are enough producers to make the gain of the speculators positive on average.

As illustrated by Figs. 3 and 4, if the number of speculators changes and that of producers is fixed the behavior is qualitatively the inverse of that of Fig. 2: The gain

of producers increases as the number of producers grows; similarly, the gain of the speculators decreases when N increases for sufficiently large N . If there are not enough producers, the game is always negative sum for the speculators, and their gain has a maximum (see Fig. 3).

We now expose exact analytical results concerning the gain of the two types of agents. They rely on the generalization of the approach of Refs. [6,7]: the calculus is carried out in detail in Appendix C. Let us introduce G_{spec} , the total gain of the speculators and G_{prod} , the one of the producers. From Eq. (15)

$$G_{\text{spec}} + G_{\text{prod}} = -\sigma^2. \quad (22)$$

The results depend on the ratio ρ between the number of producers, on the number of speculators and on c , the parameter introduced in the previous section. We obtain

$$\frac{\sigma^2}{N} = \frac{c + \rho + (1 - c)Q}{(1 + \chi)^2} + (1 - c)(1 - Q), \quad (23)$$

where χ is the magnetic susceptibility of the system, and Q is defined in Section 3.2. These two quantities depend on α and on $(1 + \rho)/(1 - c)$ (see Appendix C). The average gain per producer is

$$\frac{G_{\text{prod}}}{\rho N} = -\frac{1}{1 + \chi} \quad (24)$$

and the average gain per speculator is

$$\frac{G_{\text{spec}}}{N} = -\frac{c + \rho + (1 - c)Q}{(1 + \chi)^2} - (1 - c)(1 - Q) + \frac{\rho}{1 + \chi}. \quad (25)$$

Figs. 2–4 completely agree with analytical results; note that the small deviations are finite size effects. The fact that the gains of producers and speculators only depend on the ratio ρ and not on how many producers and speculators there are in the game explains why Figs. 3 and 4 look very much like the inverse of Fig. 2.

As it emerges for the replica calculus, the critical point α_c only depends⁵ on $(1 + \rho)/(1 - c)$ (see Fig. 5), that is, on the distribution of the quenched disorder. Numerical data (circles) completely agree with our results. The vertical line corresponds to the standard MG ($\rho = 0$ and $c = \frac{1}{2}$). A more intuitive version of this phase diagram is shown in Fig. 6 for $c = 0$.

The game becomes favorable, on average, for the speculators when their average gain is greater than zero. Using Eq. (25), one can plot the curve of zero sum gain for the speculators (see Fig. 6). One can see that the number of producers must be greater than $1.868 \dots P$ (this value depends on c) in order to make the game positive sum for the speculators; this is consistent with numerical simulations (Figs. 3 and 4).

The main message of these results is that producers always benefit from the presence of speculators, and reversely: both types of agents live in symbiosis. Indeed, the

⁵ This explains why evolutionary schemes that preserve the distribution of the quenched disorder have the same α_c [17], while others that involve Darwinism, shift α_c [1,9].

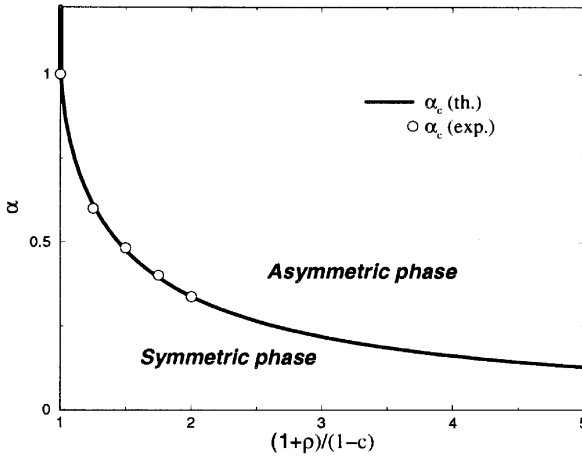


Fig. 5. Phase diagram $\alpha_c[(1 + \rho)/(1 - c)]$.

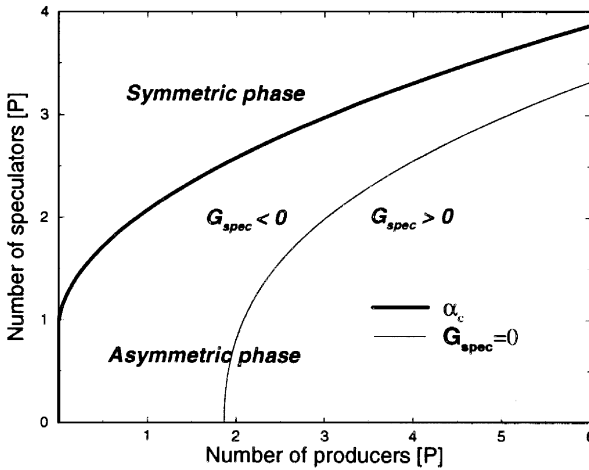


Fig. 6. Phase diagram, and zero sum game for speculators with $c = 0$ at a fixed P .

producers introduce systematic biases into the market, and without speculators, their losses would be proportional to these biases. The speculators precisely try to remove this kind of bias, reducing also systematic fluctuations in the market, thus reducing the losses of the producers and their own losses. Moreover, the efforts of speculators yield a positive gain only if the number of producers is sufficiently large. In this respect the symmetric phase, where producers do not lose and speculators lose a lot, is unrealistic: real speculators would rather withdraw from a market which is in this phase, thus increasing α , and recovering the asymmetric phase. This suggests that a grand-canonical MG is much more realistic.⁶ Here we briefly present an over-simplified

⁶ See also [18,19].




Fig. 7. Average gain per agent versus the number of producers (in P units) in the grand canonical MG ($N = 107$, $M = 5$, $\alpha = 0.3$, $S = 2$, $c = \frac{1}{2}$, average over 500 realizations).




Fig. 8. Average number of speculators versus the number of producers (in P units) in the grand canonical MG ($N = 107$, $M = 5$, $\alpha = 0.3$, $S = 2$, $c = \frac{1}{2}$, average over 500 realizations).

“grand-canonical” MG. An agent enters into the market only when she has a strategy with virtual points greater than zero. As a result, the game is always in the asymmetric phase, but almost at the transition point: the average losses of the producers are always extremely small (see Fig. 7). When the number of producers increases, the a priori asymmetry of the outcome increases, and more and more agents actually play the game (see Fig. 8), thus in this situation, the producers give incentives to play to the speculators. Accordingly, the average gain of the speculators, is much higher in this grand-canonical MG than in the corresponding canonical MG.




Fig. 9. Normalized variance of the outcome with (opaque circles and without (black squares) noise traders; the dotted line is the naive theoretical prediction. Inset: difference of variances with and without noise traders ($N = 101$ speculators, 50 noise traders, average over 1000 realizations).

6. Speculators, producers and noise traders

The debate about what the noise traders do to a competitive market is not closed [20]. In the economics literature a noise trader is not very precisely defined. Sometimes they are synonym with speculators. We define noise traders in the following way: they choose their actions without any basis. Compared with speculators, who analyze carefully the market information, noise traders take action in a purely random way (see Appendix C). Noise traders may be speculators who base their action on astrology, on “fengshui”, or on some “random number generators”. Our present model allows us to evaluate the influence of noise traders on the market. They increase the market volatility σ^2 , as shown in Fig. 9 and in Appendix C. Therefore, in principle, they do harm to themselves as well as to other participants. Actually, in the linear-payoff version that we consider, the average gain of speculators and producers is not much affected by noise traders, since $\langle A_{\text{noise}} \rangle = 0$. However, it is easy to see that in the original version, where $g_i(t) = -a_i(t) \text{sign} A(t)$, payoffs are reduced by the presence of noise traders (see Appendix D).

Our numerical results of Fig. 9 also shows that deep in the symmetric phase, noise traders reduces the volatility per agent $\sigma^2/(N + N_{\text{noise}})$, when this becomes bigger than one. This is easy to understand assuming that the only effect of noise traders is to increase σ^2 by a constant equal to $N_{\text{noise}} \equiv \eta N$. Let σ_0^2/N be the volatility per agent, without noise traders ($\eta = 0$) and σ_η^2 that with noise traders. The variation in the volatility per agent in the presence of noise traders is

$$\frac{\sigma_\eta^2}{N(1 + \eta)} - \frac{\sigma_0^2}{N} \simeq \frac{\sigma_0^2 + \eta N}{N(1 + \eta)} - \frac{\sigma_0^2}{N} = \frac{1 - \sigma_0^2/N}{1 + 1/\eta}. \tag{26}$$

As illustrated by Fig. 9, numerical simulations globally confirm these conclusions, but also show that the effects of the noise traders are more pronounced than those predicted by theory.

7. Market impact

In order to quantify the impact of an agent on the market let us first consider the case of an *external* agent with S strategies: This agent does not take part in the game but just observes it from the outside. From this position, each of her strategies gives an average⁷ virtual gain

$$u_s = -\overline{a_s \langle A \rangle}, \quad s = 1, \dots, S. \tag{27}$$

Given that the strategies a_s^μ are drawn randomly, u_s are independent random variables. Since u_s is the sum of $P \gg 1$ independent variables $a_s^\mu \langle A^\mu \rangle / P$, their distribution is Gaussian with zero mean and variance

$$\text{Var}(u_s) = \frac{1}{P^2} \sum_{\mu=1}^P \text{Var}(a_s^\mu) \langle A^\mu \rangle^2 = \frac{H}{P}.$$

Clearly, one of these strategies, that with $u_{s^*} = \max_s u_s$, is superior to all others.⁸ It would be most reasonable for this agent to just stick to this strategy.

However, the same agent *inside* the game will typically use not only strategy s^* . This is because every strategy, when used, delivers a *real* gain which is reduced with respect to the virtual one by the “market impact”. Imagine the “experiment” of injecting the new agent in a MG. Then $\langle A^\mu \rangle \rightarrow \langle A^\mu \rangle + a_s^\mu$, where, in a first approximation, we neglect the reaction of other agents to the new-comer. Then the real gain of the new-comer is

$$g_s \cong -\overline{a_s \langle A \rangle} - \langle a_s a_s \rangle = u_s - 1. \tag{28}$$

The agent will then update the scores $U_s(t)$ with the real gain g_s for the strategy she uses and with the virtual one $u_{s'} = g_{s'} + 1 - \overline{a_s a_{s'}}$, for the strategies she does not use (in the following, we neglect the term $\overline{a_s a_{s'}}$). Therefore, inductive agents over-estimate the performance of the strategies they do not play. Then if strategy s is played with a frequency p_s , the virtual score increases *on average* by

$$\begin{aligned} \delta U_s &= U_s(t+1) - U_s(t) = p_s g_s + (1 - p_s)(g_s + 1) \\ &= g_s - p_s + 1 \end{aligned} \tag{29}$$

at each time step (on average). If the agent ends up playing only n out of her S strategies with some frequency $p_s > 0$, it must be that the virtual score increases δU_s are all equal for these strategies and the virtual scores of strategies not played are lower.

⁷ The average is meant over a long time here.

⁸ The distribution of u_{s^*} can be easily computed using extreme statistics. For $S \gg 1$ typically $u_{s^*} \cong \sqrt{2H \log(S)/P}$.




Fig. 10. Ratio of real H' over approximated $H' \simeq H - 2|\overline{\langle A \rangle}| + 1$ versus α ($N = 101$, average over 100 realizations).

More precisely, let $s = 1, \dots, n$ label the strategies which are played and $r = n + 1, \dots, S$ those which are not played. It must be that

$$\delta U_s = g_s - p_s + 1 = v, \quad s = 1, \dots, n, \tag{30}$$

$$\delta U_r = g_r + 1 < v, \quad r = n + 1, \dots, S. \tag{31}$$

These equations yield the number n of strategy that this agent will use. Normalization of p_s in the first equation gives the average virtual gain v of the agent, which is

$$v = \frac{1}{n} \sum_{s=1}^n g_s - \frac{1}{n} + 1. \tag{32}$$

Using $p_s = g_s + 1 - v$, we can compute the real gain of the inductive agent $g = \sum_{s=1}^n p_s g_s$.

Summarizing, we find that inductive agents mix their best strategy with less-performing ones. This is a consequence of the fact that they neglect their impact on the market.

So far, we did not take into account the reaction of other agents to the new-comer. In order to quantify this effect, let us consider a MG in the asymmetric phase, and let us add a new agent with the *best* strategy $a^\mu = -\text{sign} \langle A^\mu \rangle$. This gives us an idea of this effect in the extreme case and we expect that for a randomly drawn strategy the effect will be smaller. Neglecting the reaction of other agents, we find that the available information with the new-comer should be $H' \simeq H - 2|\overline{\langle A \rangle}| + 1$. Fig. 10 shows that the reaction of all agents is indeed negligible, except near the critical point, where H is of the order of 1.

8. Gain

In this section we show how the behavior and the gain of each agent (speculator as well as producer) depends on her microscopic constitution and on the asymmetry of the outcome $A(t)$ in the asymmetric phase. Let us denote the gain of agent i by g_i ; by definition,

$$g_i = -\overline{\langle A a_i \rangle}. \quad (33)$$

In the asymmetric phase, since the stationary state is mean field, $\langle s_i s_j \rangle = m_i m_j$. Consequently, by expanding Eq. (33) one obtains

$$\begin{aligned} g_i &= -\overline{\langle A \rangle} \omega_i - \overline{\langle A s_i \rangle} \xi_i \\ &= -\overline{\langle A \rangle} \omega_i - \overline{\langle A \rangle} \xi_i m_i - \overline{\xi_i^2} (1 - m_i^2). \end{aligned} \quad (34)$$

Remember that the stationary behavior of agent i is described by $v_i = -2\overline{\langle A \rangle} \xi_i$ (see Section 3). If an agent is non-frozen, $v_i = 0$, while $m_i = -\text{sign } v_i$ otherwise, hence the gain of a generic agent i is

$$g_i = -\overline{\langle A \rangle} \omega_i + |\overline{\langle A \rangle} \xi_i| - \overline{\xi_i^2} (1 - m_i^2). \quad (35)$$

Note that the second term of the above equation vanishes for a non-frozen agent j and therefore

$$g_j = -\overline{\langle A \rangle} \omega_j - \overline{\xi_j^2} (1 - m_j^2) \quad \text{non-frozen}. \quad (36)$$

On the other hand, the third term of Eq. (35) vanishes if agent k is frozen:

$$g_k = -\overline{\langle A \rangle} \omega_k + |\overline{\langle A \rangle} \xi_k| \quad \text{frozen}. \quad (37)$$

In Eqs. (36) and (37), the gain of each agent is expressed as her internal constitution, allowing us to interpret what does the gain of a general agent depends on. In both equations, the first term $-\overline{\langle A \rangle} \omega_i$, which represent how much the agents lose due to their bias, is on average negative, due to the impact this bias has on the market. The second term in Eq. (36) is always negative, and represents the losses due to the switching between strategies, which, as shown above, arises from the neglect of market impact. Since the probability distribution function of m_i is not Gaussian [6], this term gives rise to a non-Gaussian distribution of g_j for non-frozen agents. The average gain of the r th best agent is represented in Fig. 11.

By contrast, the term $\overline{\xi_k^2} (1 - m_k^2)$ disappears for a frozen agent because $m_k^2 = 1$. It is replaced by $|\overline{\langle A \rangle} \xi_k|$ which is always positive and which measures how well agent k exploits the available information. Therefore, in average, the frozen agents gain more than the non-frozen ones. This is clearly illustrated in Fig. 12 which also shows that Eqs. (36) and (37) are exact. Finally, a producer is of course frozen, and her gain is always lower than zero in this phase, since she has $|\overline{\langle A \rangle} \xi_k| = 0$.




Fig. 11. Bar-code structure of the r th best agent's gain ($\alpha=10$, $M=10$, $S=2$, average over 300 realizations).

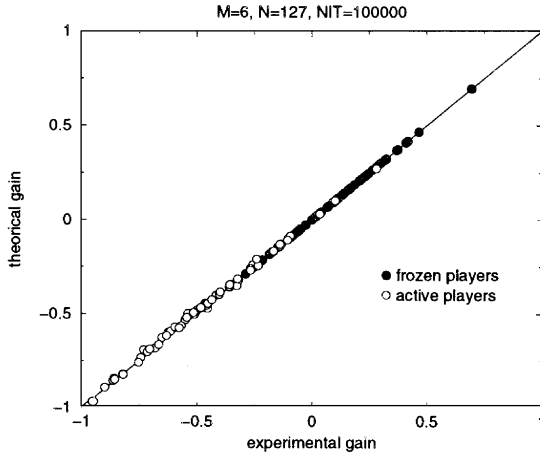


Fig. 12. Theoretical gain versus experimental gain showing that the frozen agents gain more than the active ones ($\alpha = 0.5$, $M = 6$).

9. Privileged agent or insider-trading

In this section we consider a MG where a particular agent has different characteristics. In particular, we address the question of what additional resources would be advantageous for this agent and in which circumstances. In the first subsection, we consider an agent with S' strategies (with $S' > S$, where S is the number of strategies assigned to other agents). The last two subsections are devoted to the study of effects of asymmetric information, in which an agent has access to privileged information which the other cannot access. This can be achieved in several ways. First, we consider the case of a pure population with memory M and one agent with a longer memory M' . Then we consider the case of an agent who knows, in advance, how a subset of agents plays.

9.1. An agent with S' strategies

In the symmetric phase, no matter how many strategies an agent has, there is no possibility of gaining. Therefore we focus in this section on the asymmetric phase.

As shown in Section 7, inductive agents over-estimate the performance of the strategies they do not play.

Let us consider now the case where an agent with S' strategies enters into a MG. As shown in Section 7, to a good approximation, the value of H/P is the only relevant information we need to retain of the stationary state of the MG without the special agent. This quantity encodes all other informations such as the number of producers, the number of strategies played by the agents in the MG and the value of α .

We carried out numerical simulations, and compared it to the analytical results derived in Section 7. These are shown in Fig. 13, for $H/P = 0.5$, and 14, for $H/P = 1$. The virtual gain v is always larger than the actual gain g . Even though g is less than the gain agents would get playing only their best strategy $E[g_{s^*}]$ (maximal gain), it is not much smaller and has the same leading behavior $g \propto \sqrt{\ln S}$.

Numerical simulations agree well with analytical results, apart from finite size effects which become more pronounced if H/P is small.⁹

Figs. 13 and 14 refer to values of H/P which are realistic of MG with producers. A moderately large S' suffices to obtain a positive gain $g > 0$. With $S = 2$ and without producers $H/P \sim 0.1$ at most. For these values the analytic approach suggests that, even playing only her best strategy an agent would need $S' > 750$ strategies to have a positive gain, whereas inductive agents would need more than $S' \simeq 2400$ strategies to obtain a positive gain. The same agent would find that her virtual gain becomes positive with only $S' > 8$ strategies. These results for $H/P = 0.1$ suffer from strong finite size effects (which indeed are of the order of P/H). One would need system sizes N which are well beyond what our computational resources allow to confirm these conclusions.

It is also interesting to observe that the number of strategies actually used by the inductive agent increases with S (sub-linearly) and it decreases as H/P increases (see Figs. 13 and 14). That means that if there is more exploitable information in the system, agent's behavior becomes more peaked on the best strategy.

9.2. $M' > M$

Let us consider the case of a pure population with memory M and one agent with a longer memory¹⁰ M' . Fig. 15 plots the gain of such an agent with $M' = M + 1$ as a function of α . The average gain of all agents is also shown for comparison. In the asymmetric phase the special agent receives a lower payoff, which can be understood by observing that she has a number of histories $P' = 2^{M'} = 2P$ bigger than that of the pure population. Thus her effective $\alpha' = 2\alpha$ is larger, which is detrimental in the asymmetric phase.

The gain of the special agent is the same as that of normal agents at the point where there is neither persistence, nor anti-persistence ($\alpha \simeq 0.25$ for $M = 3$, and α_c in the thermodynamic limit).

⁹This is mostly due to the term which we have neglected in Section 7: it is typically of the order of P/H .

¹⁰In this kind of numerical simulations, one has to keep the dynamics of histories.




Fig. 13. Upper graph: average number of played strategies (circles) versus S' . Below: average virtual (diamonds) and actual (squares) gains versus S' for $H/P=0.5$, from top to below (averages over 500 realizations). The lines are theoretical predictions.




Fig. 14. Upper graph: average number of played strategies (squares) versus S' . Below: average virtual (diamonds) and actual (circles) gains versus S' for $H/P=1$, from top to below (averages over 500 realizations). The lines are theoretical predictions.

In contrast, in the symmetric phase, the game is symmetric for normal agents but their anti-persistent behavior produces arbitrages who can be exploited by agents having a bigger memory. Indeed, as α decreases, the available information $H_{M'}$ for the privileged agent grows.¹¹ As a result the gain of the privileged agent becomes larger than that of other agents and as α becomes small enough, it becomes positive.

¹¹ $H_{M'}$ is defined as $H = \overline{\langle A \rangle^2}$, but with an average over $\mu' = 1, \dots, 2P$.

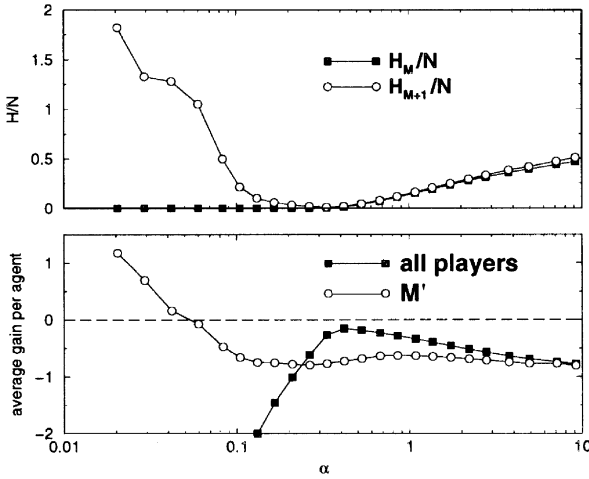


Fig. 15. Upper graph: normalized available information for M and $M + 1$. Lower graph: Gain of an agent with $M + 1$ within a pure population with $M = 3$ ($S = 2$, average over 3000 realizations).




Fig. 16. Gain of an agent with $M' = M + \Delta M$ within a pure population with $M = 3$ ($\alpha = 0.1$, average over 1000 realizations).

Can the anti-persistence be exploited even more if one increases M' ? Fig. 16 answers clearly no. This is not surprising since again the effective α is bigger and bigger as M' is increased. At the same time, the available information increases, but too slowly.

9.3. Espionage

Some agents may have access to some information about other agents. This is the case of a stock broker who knows his clients' orders before execution, hence he has privileged information and should be barred from trading. When there is no available

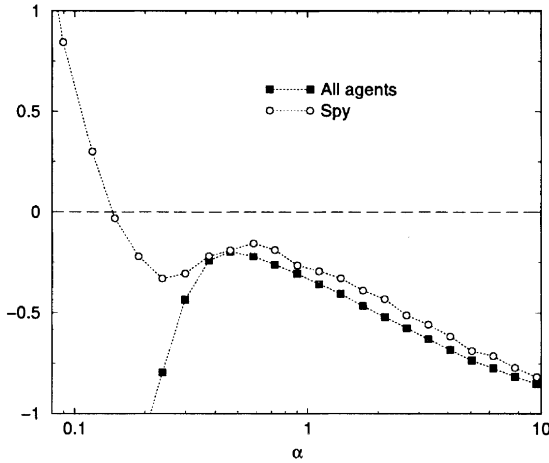


Fig. 17. Gain of a spy and average gain of all agents versus α ($N = 101$, $N_B = 3$, $100P$ iterations, average over 100 realizations).

information, as in the symmetric phase, an agent who has access to asymmetric information can expect at least to lose much less than the other agents, or even have a positive gain. Also, since having access to a little information is greatly preferable to no information at all, only a very limited amount of information is needed to get a considerable advantage. Suppose that agent b knows the sign $s_{\mathcal{B}}$ of the aggregate actions of a subset \mathcal{B} of other agents. Let $B = |\mathcal{B}|$ be the number of agents in \mathcal{B} . Then $s_{\mathcal{B}}(t) = \text{sign} \sum_{i \in \mathcal{B}} a_i(t)$. She can exploit this supplementary information by having two virtual values $U_{b,s}^+(t)$ and $U_{b,s}^-(t)$ for each of her strategies. In other words, if agent b knows that $s_{\mathcal{B}}(t) = +1$ before having to choose, she takes her decision according to the scores $U_{b,s}^+(t)$, that is,

$$s_b(t) = \arg \max_{s=1, \dots, S} U_{b,s}^+(t), \tag{38}$$

she updates the scores of her strategies according to

$$U_{b,s}^+(t+1) = U_{b,s}^+(t) - a_{b,s}^{\mu(t)} A^{\mu(t)} \tag{39}$$

and analogously if $s_{\mathcal{B}}(t) = -1$.

What is the kind of the supplementary information this agent has access to?

Since the outcome is anti-persistent in the symmetric phase and persistent in the asymmetric phase, only at the critical point there is no long-term correlation in the outcome [13]. Accordingly, the spy always gains more than the average, except at the critical point where she gains the same (see Fig. 17). With this setting, the agent has access in particular to the anti-persistence of the symmetric phase, explaining why even if only one agent is spied, the gain of the broker is much bigger (Fig. 18).

Finally, the comparison between the two types of asymmetric information we have considered shows that it is much more interesting to spy than to have a larger memory:




Fig. 18. Gain of a spy versus the number of spied agents ($N=1001$, $\alpha=0.15$, average over 1000 realizations).

in the former case, one is sure to win more than the normal agents, except at the critical point.

10. Conclusions

In this work we have shown how to ask questions about real market mechanisms in a toy model. In spite of the severe simplification of the MG, with little modification one is able to study a broad spectrum of problems which could be dreamed of previously. The central result is to show that agents with limited rationality (or limited information processing power) can only make a market marginally efficient. To the first approximation one can say that these inductive players can maintain an approximate equilibrium, which is the central result of the El-Farol model. But studying carefully the fluctuations one finds that the fact that the market is more or less efficient does not imply that one can stop playing and sit at a randomly chosen site. Doing so would make the model less efficient. It is around this residual (marginal) inefficiency that the players are busy about.

With the introduction of producers the game can be of positive sum. We have shown how producers and speculators live in a symbiosis: producers are passive players who do not try to switch strategies. The reason is that they voluntarily give up the speculation opportunities because they have outside business in mind. Thus, producers inject information that the eager speculators are just happy to feed on. The speculators, while making away profits, perform a social function by providing liquidity thus reducing producers' market impact. We believe this is also true in real markets. Numerous other results show that it is now possible to systematically study markets with heterogeneous agents, with real questions in mind.

Appendix A. Geometric and algebraic approaches to the MG

This appendix is devoted to giving intuitive but rigorous views of what happens in the MG.

A.1. Geometric approach to the MG

The global behavior of the MG, measured by σ^2 , can be quite well understood with a geometrical approach. Indeed, it is directly related to a much more intuitive geometrical concept: the Hamming distance between agents [9], which is defined as follows for agents i and j :

$$\overline{d_{i,j}} = \frac{\overline{(a_i - a_j)^2}}{4} = \frac{1}{2} - \frac{1}{2}\overline{a_i a_j}. \quad (\text{A.1})$$

It is worthwhile to note that $1 - \overline{d_{i,j}}$ equals the probability that both agents take the same action for a randomly drawn history, so for an agent, maximizing her distance with respect to all other agents is equivalent to maximizing her gain. Since the game is dynamical, one has to consider the time average of the actual Hamming distance between those two agents

$$\overline{\langle d_{i,j} \rangle} = \frac{1}{2} - \frac{1}{2}\overline{\langle a_i a_j \rangle}. \quad (\text{A.2})$$

The average Hamming distance per agent is then

$$\overline{\langle d \rangle} = \frac{1}{N(N-1)} \sum_{i,j} \overline{\langle d_{i,j} \rangle} = \frac{1}{2} - \frac{1}{2N(N-1)} \sum_{i,j} \overline{\langle a_i a_j \rangle}. \quad (\text{A.3})$$

The relationship between the distance and the fluctuations arises naturally by rewriting the latter as

$$\sigma^2 = N + \sum_{i \neq j} \overline{\langle a_i a_j \rangle}, \quad (\text{A.4})$$

that is, as a sum over random fluctuations and correlations. Putting Eqs. (A.3) and (A.4) together, one finds

$$\frac{\sigma^2}{N} = 1 - 2(N-1) \left(\frac{1}{2} - \overline{\langle d \rangle} \right). \quad (\text{A.5})$$

This equation¹² links the geometrical [9,12] and the analytical approaches [6,13]. It states that finding the average Hamming distance between the agents is equivalent to determining σ^2 by the analytical tools used in [6,7,13]. In general, it is impossible to find the average distance with a geometrical approach due to the fact that the Hamming distance is not transitive.¹³ However, in the so-called reduced space of strategies (RSS) [9], the distance is transitive, consequently Johnson et al. could find an approximate

¹² It is exact for any S ; even more, it remains exact if agents do not have the same number of strategies.

¹³ The knowledge of $d_{i,j}$ and $d_{i,k}$ does not allow that of $d_{j,k}$.

analytical expression of $\overline{\langle d \rangle}$, and, by implicitly using Eq. (A.5) (which is straightforward in the RSS), they also gave an approximative expression of σ^2 [21]. An equation quite similar to Eq. (A.5) also appears in [22], where it is shown that perceptrons playing the MG can cooperate.

A.2. Algebraic approach to the phase transition

We expose the algebraic origin of the phase transition. As it has been recalled, the agents actually try to minimize the available information H [6,7], and can actually cancel it when $\alpha < \alpha_c$. Let us see why. Since H is a sum of P non negative averages $\overline{\langle A \rangle^2}$, $H = 0$ only if all averages are zero, namely $\overline{\langle A \rangle} = 0 \forall \mu$, or equivalently

$$\sum_{i=1}^N \xi_i^\mu \langle s_i \rangle = -\Omega^\mu \quad \forall \mu. \quad (\text{A.6})$$

These are P linear equations in N variables. However, the N variables $m_i = \langle s_i \rangle$ are restricted to the $[-1, 1]$ interval. Above α_c there are $N\phi$ variables which are frozen at the boundary of this interval ($m_i = \pm 1$). Therefore there are $(1 - \phi)N$ free variables only. As shown in Refs. [6,7], the point α_c marks the transition below which the system of equations (A.6) becomes degenerate, i.e., when there are more variables than equations. Exactly at α_c the number of free variables $(1 - \phi)N$ exactly matches the number of equations P . Dividing this equation by N gives an equation for α_c ,

$$\alpha_c = 1 - \phi \quad (\text{A.7})$$

which is indeed confirmed numerically to a high accuracy.

When $\alpha < \alpha_c$, there are more free variables (N indeed) than equations: the solutions of Eq. (A.6) then belong to a subspace of dimension $N - P$. This allows the anti-persistent behavior to take place, because the system is free to move on this subspace. In the special case $c = 0$, since there is no bias $\Omega^\mu = 0$. The linear system of equation is then homogeneous and the solution $\langle s_i \rangle = 0$ for all i always exists. In particular, if $\alpha > 1$, this solution is unique, hence $\sigma^2/N = 1$. When $\alpha < 1$, a subspace of solutions of dimension $N - P$ arises, and the anti-persistent behavior also takes place. Note that in this case, the system is always in the symmetric phase, therefore there is no phase transition.

This argument easily generalizes to $S > 2$ strategies [7]. If agents use, on average, $n(S)$ strategies (and $S - n(S)$ are never used) the number of free variables is $Nn(S)$. There are P plus N equations which these have to satisfy, where the latter N comes from the normalization condition on the frequency with which each strategy is used. At the critical point, these two numbers are equal, and we find

$$n_c(S) = \alpha_c(S) + 1. \quad (\text{A.8})$$

At the critical point nearly one half of the strategies yield positive virtual gain and are used, whereas the others are not used [7]. From this we find

$$\alpha_c(S) \cong \alpha_c(2) + \frac{S}{2} - 1. \quad (\text{A.9})$$

This shows that actually α_c grows linearly with S , but in a slightly less-simple way than that previously believed [2,9,13,21].

Let us now show how the behavior of the agents is related to persistence/anti-persistence. We define W as the average over the agents of v_i^2 :

$$W = \frac{1}{N} \sum_{i=1}^N (v_i)^2 \tag{A.10}$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T^2} \sum_{t,t'=1}^T \sum_{i=1}^N \frac{1}{N} \xi_i^{\mu(t)} \xi_i^{\mu(t')} A(t)A(t'), \tag{A.11}$$

since the ξ_i^μ are independently drawn,

$$\frac{1}{N} \sum_{i=1}^N \xi_i^{\mu(t)} \xi_i^{\mu(t')} = (1 - c)\delta_{\mu(t),\mu(t')} + O(1/\sqrt{N}). \tag{A.12}$$

Thus for large N

$$W \simeq \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \overline{\langle A(t)A(t - \tau) | \mu(t) = \mu(t - \tau) \rangle}, \tag{A.13}$$

where $\overline{\langle A(t)A(t - \tau) | \mu(t) = \mu(t - \tau) \rangle}$ means that the average is taken over time $\tau = 0, \dots, t - 1$ for all t and τ such that $\mu_t = \mu_{t-\tau}$, and summed over all histories. A closely related quantity was first studied in Ref. [13] where it was shown to quantify anti-persistence in the symmetric phase.

Note that this equation implies that there can be no frozen agents unless the outcome exhibits persistence (i.e., $\overline{\langle A(t)A(t - \tau) | \mu(t) = \mu(t - \tau) \rangle} \neq 0$), which agrees with the analysis of [13]. In this case we find

$$W \simeq \frac{H}{N}. \tag{A.14}$$

Furthermore, the condition of freezing $v_i \neq 0$ is equivalent to [13]

$$\overline{\xi_i \xi_i} < |\tilde{h}_i|, \tag{A.15}$$

where $\tilde{h}_i = \overline{\Omega \xi_i} + \sum_{j \neq i} \overline{\xi_i \xi_j} \langle s_j \rangle$. It is worthwhile to see that $\overline{\xi_i \xi_i}$ is the internal hamming distance. Eqs. (A.13) and (A.14) give global conditions whether there can be frozen players or not, while Eq. (A.15) give conditions on individual freezing.

Appendix B. The MG in biology

The MG model has another important application in biology: the sex ratio of 50 : 50. In the widely read book of Richard Dawkins “Selfish Gene” [23], the Fischer theory was brilliantly explained: if in the offspring pool either males or females were in minority, reproductive strategies for giving birth to a member in that minority would enjoy a genetic advantage linearly proportional to the deviation from the 50 : 50 ratio.

The stable ratio is thus dynamically maintained. Brian Arthur's "El Farol" model, is also of the same genre, to show that using alternative strategies can lead to equilibrium. MG goes one step further: while the equilibrium point is previously solved in different contexts by Fisher, Arthur et al., we concentrate on more refined questions.

Appendix C. Replica method for the MG

For the sake of generality, we consider three different population of agents:

(1) The first population is composed of N *speculators*. These are adaptive agents and they have each two *speculative* strategies $a_{\uparrow,i}^{\mu}, a_{\downarrow,i}^{\mu}$ for $i = 1, \dots, N$ and $\mu = 1, \dots, P$. These are drawn at random from the pool of all strategies, independently for each agent. We allow a correlation among the two strategies of the same agent:

$$P(a_{\uparrow}, a_{\downarrow}) = \frac{c}{2} [\delta_{a_{\uparrow},+1} \delta_{a_{\downarrow},+1} + \delta_{a_{\uparrow},-1} \delta_{a_{\downarrow},-1}] + \frac{1-c}{2} [\delta_{a_{\uparrow},-1} \delta_{a_{\downarrow},+1} + \delta_{a_{\uparrow},+1} \delta_{a_{\downarrow},-1}]. \quad (C.1)$$

Note that, for $c=0$ agents choose just one strategy a_{\uparrow} and fix $a_{\downarrow} = -a_{\uparrow}$ as its opposite, whereas for $c=1$ they have one and the same strategy $a_{\uparrow} = a_{\downarrow}$. The original random case [1,8] corresponds to $c = \frac{1}{2}$. These agents assign scores $U_{s,i}(t)$ to each of their strategies and play the strategy $s_i(t)$ with the highest scores, as discussed in the text. Therefore for speculators

$$a_{\text{spec}}(t) = a_{s_i(t),i}^{\mu(t)}. \quad (C.2)$$

(2) Then we consider $N_{\text{prod}}^{\text{indep}} = \rho N$ *producers*: They have only one randomly and independently drawn strategy b_i^{μ} so

$$a_{\text{prod}}(t) = b_i^{\mu(t)}. \quad (C.3)$$

Producers have a predictable behavior in the market and they are not adaptive. Instead of ρN *independent* producers one can also consider $N_{\text{prod}}^{\text{dep}}$ correlated producers who all have the same predictable behavior b_{prod}^{μ} .

(3) Finally, we consider ηN *noise traders*. These are defined as agents whose actions are given by

$$a_{\text{noise}}(t) = \text{random sign}. \quad (C.4)$$

Each noise trader as a random number generator which is independent of those other agents.

It has been shown [6,7] that the stationary state properties of the MG are described by the ground state of H . Note that this approach fails however to reproduce the anti-persistent behavior which is at the origin of crowd effects in the symmetric phase. In our case

$$A(t) = A_{\text{spec}}(t) + A_{\text{prod}}(t) + A_{\text{noise}}(t), \quad (C.5)$$

where

$$A_{\text{spec}}(t) = \sum_{j=1}^N a_{s_j(t),j}^{\mu(t)} \tag{C.6}$$

and

$$A_{\text{prod}}(t) = \sum_{j=1}^{\rho N} b_j^{\mu(t)} \equiv A_{\text{prod}}^{\mu(t)} \tag{C.7}$$

and $A_{\text{noise}}(t) = 2k(t) - \eta N$, where $k(t)$ is a binomial random variable with $P(k) = \binom{\eta N}{k} 2^{-\eta N}$. Since $H = \overline{\langle A \rangle^2}$ and the contribution of noise traders to $\langle A^\mu \rangle$ vanishes $\langle A_{\text{noise}} \rangle = 0$, the collective behavior of the system is independent of η . Noise traders shall contribute a constant ηN to σ^2 and will not affect other agents. This only holds in the asymmetric phase (see text). We can then reduce to the study of speculators and producers only.

Let us define, for convenience, $A^\mu = A_{\text{spec}}^\mu + \lambda A_{\text{prod}}^\mu$, where

$$A_{\text{spec}}^\mu = \sum_{i=1}^N \left[a_{\uparrow,i}^\mu \frac{1 + s_i}{2} + a_{\downarrow,i}^\mu \frac{1 - s_i}{2} \right] \tag{C.8}$$

and A_{prod}^μ is given in Eq. (C.7). Here s_i is the dynamical variable controlled by speculator i . We shall implicitly consider directly time-averaged quantities, so s_i is a real variable in $[-1, 1]$ rather than a discrete one. The parameter λ is inserted so that, once we have computed the energy $H = \overline{(A_{\text{spec}} + \lambda A_{\text{prod}})^2}$ we can compute the total gain G_{prod} of producers by

$$G_{\text{prod}} \equiv -\overline{AA_{\text{prod}}} = - \left. \frac{1}{2} \frac{\partial H}{\partial \lambda} \right|_{\lambda=1} .$$

The gain of speculators is obtained subtracting this contribution and that of noise traders from the total gain $-\sigma^2$

$$G_{\text{spec}} = -\sigma^2 + \eta N - G_{\text{prod}} . \tag{C.9}$$

C.1. Replica calculation

The zero temperature behavior of the Hamiltonian H can be studied with spin-glass techniques [15,14]. We introduce n replicas of the system, each with dynamical variables $s_{i,c}$, labeled by replica indices c , $d = 1, \dots, n$. Then we write replicated partition function

$$\langle Z^n(\beta) \rangle = \text{Tr}_s \prod_{\mu,c} \langle e^{-\beta/P(A_c^\mu)^2} \rangle_{a,b} , \tag{C.10}$$

where the average is over the disorder variables $a_{s,i}^\mu$, b_i^μ and Tr_s is the trace on the variables $s_{i,c}$ for all i and c . Following standard procedures [15,14], we introduce a Gaussian variable z_c^μ so that we can linearize the exponent in Eq. (C.10). This allows us

to carry out the averages over a 's and b 's explicitly. Then we introduce new variables $Q_{c,d}$ and $r_{c,d}$ with the identity

$$1 = \int dQ_{c,d} \delta \left(Q_{c,d} - \frac{1}{N} \sum_i s_{i,c} s_{i,d} \right) \\ \propto \int dr_{c,d} dQ_{c,d} e^{-\alpha\beta^2/2r_{c,d}(NQ_{c,d} - \sum_i s_{i,c} s_{i,d})}$$

for all $c \geq d$, which allow us to write the partition function (to leading order in N) as

$$\langle Z^n(\beta) \rangle = \int d\hat{Q} d\hat{r} e^{-Nn\beta F(\hat{Q}, \hat{r})}$$

with

$$F(\hat{Q}, \hat{r}) = \frac{\alpha}{2n\beta} \text{Tr} \log \hat{T} + \frac{\alpha\beta}{2n} \sum_{c \leq d} r_{c,d} Q_{c,d} - \frac{1}{n\beta} \log [\text{Tr}_s e^{\alpha\beta^2/2 \sum_{c \leq d} r_{c,d} s_c s_d}]. \tag{C.11}$$

The matrix \hat{T} is given by

$$T_{a,b} = \delta_{a,b} + \frac{2\beta}{\alpha} [c + \rho + (1 - c)Q_{a,b}].$$

For *correlated* producers we would have obtained the same result but with $\rho \rightarrow \rho + \rho^2 N \varepsilon^2$, where ε measures the bias of producers towards a particular action for a given μ , or equivalently the correlation between the actions of two distinct producers. More precisely, ε^2 is the average of $b_i^\mu b_j^\mu$ for $i \neq j$ and for all μ . Therefore, the limit $\rho \rightarrow \infty$ also corresponds to a small share of producers $\rho \ll 1$ with a small bias $\varepsilon \neq 0$. Note that a bias $\varepsilon \sim \sqrt{N}$ corresponds indeed to $\sim N$ independent producers. Equivalently $\sim \sqrt{N}$ correlated producers, with ε finite are equivalent to $\sim N$ independent producers.

With the replica symmetric ansatz

$$Q_{c,d} = q + (Q - q)\delta_{c,d}, \quad r_{c,d} = 2r + (R - 2r)\delta_{c,d}$$

the matrix \hat{T} has $n - 1$ degenerated eigenvalues $\lambda_0 = 1 + 2(1 - c)\beta(1 - q)/\alpha$ and one eigenvalue equal to $\lambda_1 = 2\beta[c + \rho + (1 - c)q]/\alpha n + 1 + 2(1 - c)\beta(1 - q)/\alpha$ therefore, after standard algebra,

$$F^{(RS)}(q, r) = \frac{\alpha}{2\beta} \log \left[1 + \frac{2(1 - c)\beta(Q - q)}{\alpha} \right] + \frac{\alpha[c + \rho + (1 - c)q]}{\alpha + 2(1 - c)\beta(Q - q)} \\ + \frac{\alpha\beta}{2} (RQ - rq) - \frac{1}{\beta} \left\langle \log \int_{-1}^1 ds e^{-\beta V_z(s)} \right\rangle, \tag{C.12}$$

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

$$V_z(s) = -\frac{\alpha\beta(R - r)}{2} s^2 - \sqrt{\alpha r} z s \tag{C.13}$$

so that the last term of $F^{(RS)}$ looks like the free energy of a particle in the interval $[-1, 1]$ with potential $V_z(s)$ where z plays the role of disorder.

The saddle point equations are given by

$$\frac{\partial F^{(RS)}}{\partial q} = 0 \Rightarrow r = \frac{4(1-c)[c + \rho + (1-c)q]}{[\alpha + 2(1-c)\beta(Q-q)]^2}, \tag{C.14}$$

$$\frac{\partial F^{(RS)}}{\partial Q} = 0 \Rightarrow \beta(R-r) = -\frac{2(1-c)}{\alpha + 2(1-c)\beta(Q-q)}, \tag{C.15}$$

$$\frac{\partial F^{(RS)}}{\partial R} = 0 \Rightarrow Q = \langle\langle s^2 \rangle\rangle, \tag{C.16}$$

$$\frac{\partial F^{(RS)}}{\partial r} = 0 \Rightarrow \beta(Q-q) = \frac{\langle\langle sz \rangle\rangle}{\sqrt{\alpha r}}, \tag{C.17}$$

where $\langle\langle \cdot \rangle\rangle$ stands for a thermal average over the above-mentioned one-particle system.

In the limit $\beta \rightarrow \infty$ we can look for a solution with $q \rightarrow Q$ and $r \rightarrow R$. It is convenient to define

$$\chi = \frac{2(1-c)\beta(Q-q)}{\alpha} \quad \text{and} \quad \zeta = -\sqrt{\frac{\alpha}{r}}\beta(R-r) \tag{C.18}$$

and to require that they stay finite in the limit $\beta \rightarrow \infty$. The averages are easily evaluated since, in this case, they are dominated by the minimum of the potential $V_z(s) = \sqrt{\alpha r}(\zeta s^2/2 - zs)$ for $s \in [-1, 1]$. The minimum is at $s = -1$ for $z \leq -\zeta$ and at $s = +1$ for $z \geq \zeta$. For $-\zeta < z < \zeta$ the minimum is at $s = z/\zeta$. With this we find

$$\langle\langle sz \rangle\rangle = \frac{1}{\zeta} \operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right) \tag{C.19}$$

and

$$\langle\langle s^2 \rangle\rangle = Q = 1 - \sqrt{\frac{2}{\pi}} \frac{e^{-\zeta^2/2}}{\zeta} - \left(1 - \frac{1}{\zeta^2}\right) \operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right). \tag{C.20}$$

With some more algebra, one easily finds

$$\chi = \left[\alpha / \operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right) - 1 \right]^{-1}. \tag{C.21}$$

Finally, ζ is fixed as a function of α by the equation

$$\sqrt{\frac{2}{\pi}} \frac{e^{-\zeta^2/2}}{\zeta} + \left(1 - \frac{1}{\zeta^2}\right) \operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right) + \frac{\alpha}{\zeta^2} = \frac{1 + \rho}{1 - c}. \tag{C.22}$$

Note that ζ only depends on the combination $(1 + \rho)/(1 - c)$ which runs from $1 -$ for $\rho = c = 0$, i.e., no producers and “perfect” speculators – to ∞ . The latter limit occurs either if $c \rightarrow 1$, i.e., when speculators become producers, or if $\rho \rightarrow \infty$ (many producers).

Eq. (C.21) means that χ diverges when $\alpha \rightarrow \alpha_c(\rho, c)^+$, which then implies that at the critical point

$$\operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right) = \alpha = \alpha_c. \tag{C.23}$$

Substituting this in the other saddle point equations, yields the following equation for $\zeta = \zeta_c$:

$$\sqrt{\frac{2}{\pi}} \frac{e^{-\zeta_c^2/2}}{\zeta_c} + \operatorname{erf}\left(\frac{\zeta_c}{\sqrt{2}}\right) = \frac{1 + \rho}{1 - c}. \tag{C.24}$$

The free energy, at the saddle point, for $\beta \rightarrow \infty$, is

$$F^{(RS)} = \frac{c + (1 - c)Q + \rho}{(1 + \chi)^2}, \tag{C.25}$$

where Q and χ take their saddle point values, Eqs. (C.20) and (C.21).

The gain of producers, from Eq. (C.12), is

$$\frac{G_{\text{prod}}}{N} = -\frac{\rho}{1 + \chi} \tag{C.26}$$

and that of speculators is obtained from Eq. (C.9).

At α_c $\chi \rightarrow \infty$ so that $F^{(RS)} \rightarrow 0$. Note that the loss of producers vanishes $L_{\text{prod}} \rightarrow 0$ as $\alpha \rightarrow \alpha_c$, whereas the loss of speculators $L_{\text{spec}} = (1 - Q)/2$ is always positive below α_c .

The phase diagram is shown in Fig. 5. here we discuss some limits.

Appendix D. The sign MG

The original MG [1] is defined with payoffs

$$g_i = -a_i(t)\operatorname{sign} A(t) \tag{D.1}$$

Over a long period of time T , the change in $\Delta_i(t)$ is given by

$$\begin{aligned} \frac{\Delta_i(t + T) - \Delta_i(t)}{T} &= \frac{-2}{T} \sum_{\tau=t}^{T-1} \xi_i^{\mu(\tau)} \operatorname{sign} A(\tau) \\ &\simeq \frac{-2}{P} \sum_{\mu=1}^P [2\operatorname{Prob}\{A(\tau) > 0 | \mu(\tau) = \mu\} - 1] \xi_i^\mu. \end{aligned} \tag{D.2}$$

Then, for any fixed μ , the relevant quantity is the probability that $A(t) > 0$, when $\mu(t) = \mu$. This can be computed within our mean-field approximation: Indeed if $\langle s_i \rangle = m_i$ we can regard s_i as a random variable with distribution

$$P(s_i = \pm 1) = \frac{1 \pm m_i}{2}.$$

Then, the relation $A^{\mu(t)}(t) = \Omega^\mu + \sum_{i=1}^N \xi_i^\mu s_i$ implies that we can consider A^μ as a Gaussian variable with variance

$$\operatorname{Var}(A^\mu) = \sum_{i=1}^N \xi_i^{\mu 2} (1 - m_i^2) + \eta N.$$

This allows us to compute

$$\text{Prob}\{A(\tau) > 0 | \mu(\tau) = \mu\} = \frac{1}{2} \text{erfc} \left(\frac{\langle A^\mu \rangle}{\sqrt{2 \text{Var}(A^\mu)}} \right).$$

Note that close to the critical point α_c , $\langle A^\mu \rangle$ is very small compared to $\text{Var}(A^\mu)$, which means that it is legitimate to expand the erfc function to a linear order. This gives us back a linear minority game, but with

$$g_i = - \frac{a_i(t)A(t)}{\sqrt{2 \text{Var}(A^\mu)}}. \quad (\text{D.3})$$

Note, then that when η increases the gains for each speculator decreases. This is actually true even away from α_c . It is indeed easy to check that $\langle \text{sign} A(t) \rangle$ decreases as η increases.

References

- [1] D. Challet, Y.-C. Zhang, *Physica A* 246 (1997) 407 (adap-org/9708006).
- [2] Y.-C. Zhang, *Europhys. News* 29 (1998) 51 (cond-mat/9803308).
- [3] H. Simon, *Models of Bounded Rationality*, MIT Press, Cambridge, 1997.
- [4] W.B. Arthur, *Amer. Econ. Assoc. Papers Proc.* 84 (1994) 406 (<http://www.santafe.edu/arthur/Papers/ElFarol.html>).
- [5] See the Minority Game's web page on <http://www.unifr.ch/econophysics>.
- [6] D. Challet, M. Marsili, R. Zecchina, preprint, cond-mat/9904392.
- [7] M. Marsili, D. Challet, R. Zecchina, preprint, cond-mat/9908480.
- [8] R. Savit, R. Manuca, R. Riolo, *Phys. Rep. Lett.* 82(10) (1999) 2203 (adap-org/9712006).
- [9] D. Challet, Y.-C. Zhang, *Physica A* 256 (1998) 514 (cond-mat/9805084).
- [10] A. Cavagna, preprint, cond-mat/9812215, 1998.
- [11] R.D. Luce, *Individual Choice Behavior: A Theoretical Analysis*, New York, Wiley, 1959.
- [12] A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, preprint, cond-mat/9903415, 1999.
- [13] D. Challet, M. Marsili, preprint, cond-mat/9904071.
- [14] M. Mezard, G. Parisi, M.A. Virasoro, *Spin Glass Theory and Beyond*, World Scientific, Singapore, 1987.
- [15] V. Dotsenko, *An Introduction to the Theory of Spin Glasses and Neural Networks*; World Scientific Publishing, Singapore, 1995.
- [16] Y.-C. Zhang, *Physica A* 269 (1999) 30 (cond-mat/9901243).
- [17] Y. Li et al., preprints, cond-mat/9903415 and cond-mat/9906001.
- [18] M. Marsili, R. Zecchina, forthcoming.
- [19] N.F. Johnson et al., forthcoming.
- [20] J.B. Delong, A. Schleifer, L.H. Summers, R.J. Waldmann, *Noise trader risk in financial markets*, *J. Political Econ.* 98 (1990) 703.
- [21] N.F. Johnson et al., preprint, cond-mat/9811227.
- [22] W. Kinzel, R. Metzler, I. Kanter, preprint, cond-mat/9906058.
- [23] R. Dawkins, *The Selfish Gene*, Oxford Univ. Press, New York, 1976.