Statistical physics of interacting agents models: Minority Games

Matteo Marsili

April 16, 2004

Contents

1 Introduction

The availability of massive amounts of financial data provides scientists the access to microscopic behavior down to individual interactions. Such a wealth of empirical information, which in principle makes it possible to test theoretical insights on financial markets to an unprecedented precision in socio-economic systems, has been one of the causes of the recent interest of physicists in financial markets. One of the contributions of physicists has been an empirical approach to financial fluctuations [Mantegna and Stanley, 2000, Bouchaud and Potters, 2000] independent of the econometric approach and often in contrast with the axiomatic approach of theoretical finance [Farmer, 1999]. The empirical evidence depicts financial markets as complex self-organizing critical systems: The statistics of real market returns deviate considerably from Gaussian statistics. Market returns display fat tails, scaling and long range volatility correlations [Mantegna and Stanley, 2000, Bouchaud and Potters, 2000]. These anomalous properties of financial fluctuations have been called the stylized facts. Agent based models of financial markets have shown that these empirical features can be self-generated by a system of more or less rational traders, interacting through a market mechanism [Farmer and Joshi, 99, Lux and Marchesi, 1999, Cont and Bouchaud, 2000, Arthur et al., 1997, Caldarelli et al., 1997, Levy et al., 2000]. A modeling approach in the spirit of statistical mechanics seems particularly appropriate, for financial markets, because features such as the stylized facts evoke the theory of critical phenomena. This explains how and when anomalous fluctuations do emerge from the interaction of many microscopic degrees of freedom. Unfortunately, for most agent based models, progress beyond the mere numerical simulation approach is hard, because of the model's complexity. So the relation with critical phenomena has not been pushed further. The Minority Game, though being a very crude representation of a financial market, has filled this gap.

The Minority Game[Challet and Zhang, 1997, Challet,] (MG) was initially designed as the most drastic simplification of Arthur's famous El Farol's Bar problem [Arthur, 1994]: it describes a system where many heterogeneous agents interact through a price system they all contribute to determine. The MG is an highly stylized model of such a situation: it captures some key features of a generic market mechanism and the basic interaction between agents and public information – i.e. how agents react to information and how these reactions modify the information itself. In addition, it allows to study in details how macroscopic quantities depend on microscopic behaviors.

However, the basic MG is a so stylized model of a financial market that prices are not even explicitly defined. Furthermore the micro-economic behavior of agents is quite simplified: agents have heterogeneous strategies but they enter the game with the same weight. In other words there are not poorer or richer agents and their wealth does not change according to their performance. Also all agents are constrained to play, with the same frequency, no matter how much they may loose. All these unrealistic features makes it hard to accept the MG as a model of a real financial market, especially when compared to other agent-based models [Lux and Marchesi, 1999, Caldarelli et al., 1997, Levy et al., 2000, Cont and Bouchaud, 2000] which have so far been more successful in reproducing the stylized facts of high frequency statistics of prices [Mantegna and Stanley, 2000].

The same stylized nature of the MG however, allows one to gain a deep understanding of its extremely rich collective behavior: Statistical mechanics of disordered systems indeed allows for a full analytic solution in the limit of infinitely many agents [Marsili et al., 2000]. More precisely, these techniques allows one to fully characterize the evolutionary equilibrium of the dynamical learning process in a truly complex system of interacting adaptive agents. In top-down approach to real financial markets, where complexity is added in steps, the analytic solution of the MG provides an invaluable starting point which allows us to keep full control on the emergent features.

The Minority Game not only captures the essential interaction between traders in a market but also provides a very simple description of market's behavior in terms of two key quantities, predictability and volatility. Examining the interplay between market's informational inefficiency (i.e. predictability) and speculative behavior, one finds very good theoretical reasons to expect that markets operate rather close to criticality [Zhang, 1999]. These expectations are fully supported by microscopic agent based market models based on the Minority Game: The picture offered by these synthetic markets is one where speculation drives market to information efficiency, i.e. to a point where market returns are unpredictable. But the point where markets become exactly efficient is the locus of a phase transition. Close to the phase transition the behavior of synthetic markets is characterized by the observed stylized facts (fat tails and long range correlations) whereas far from the critical region the market is well described in terms of random walks.

1.1 A plan of these notes

A good start on beginning to understand the Minority Game is to program it. Since the computer code produces all the complexity of the Minority Game, the first step is to read the key dynamical equations from it. The second step is to simplify further the model so as to leave unaffected its generic behaviour, making an analytic treatment possible. Identifying unnecessary complications is not an easy task. For example, one *a priori* essential fact that a model of a financial market should capture is that traders base their strategies on processing the information about past price moves which are produced by their very actions. Very few would have guessed that such a feedback loop of information *a posteriori* turns out to have weak consequences on the collective behaviour [Cavagna, 1999]. Some other 'simplifications' actually amount in making the model apparently more complex, but remove hindrances to a mathematical approach. It is the case of the 'temperature' introduced in Ref. [Cavagna et al., 1999] in order to model agents who choose in a probabilistic fashion. The Minority Game has reacted in a totally counter-intuitive way to this modification, as we shall see.

After having defined the model and discussed how it can be simplified, we shall discuss how the Minority Game can be derived from a market mechanism. Then we shall outline the generic behavior of the Minority Game.

The simplifications discussed above bring the Minority Game in the range of applicability of the powerful tools of statistical physics. First, by taking the average of dynamical equations, it is possible to characterize the "average behaviour" in the stationary state. The problem can be cast into the form of the minimization of a function H with respect to the "average behaviour" of agents[1]. Since H quantifies market predictability, in loose words we can say that agents in the Minority Game make the market as unpredictable as possible, given their strategies. This is in line with *no-arbitrage* arguments: if an agent is able to predict the market shewould modify herbehaviour in order to exploit this information. Therefore, in the stationary state, the agents either have no possibility to predict the market or already exploit the information they have.

The minimization of H reveals that the collective behavior of the Minority Game features a *phase transition* as a function of the number N of agents. When there are less agents than a critical number, the price evolution seems predictable to an external agent (but not to those already playing), whereas when the number of agents is beyond the critical number, the market becomes unpredictable. This suggests that, as long as there are few participants, the market will attract more and more agents, thus approaching the critical number where the market becomes unpredictable and hence unattractive. Hence the scenario depicted by the Minority Game lends support to the conjecture that markets *self-organize at a critical point*, an issue we leave for the next chapter. Apart from predictability, the two phases have quite different properties: in the predictable phase (when there are few agents) the stationary state is unique and ergodic, i.e. independent of initial conditions. When there are many agents (unpredictable phase) the stationary state is not ergodic as it remembers (i.e. depends on) the initial conditions. In particular the fluctuations in the stationary state is a decreasing function of the degree of heterogeneity of initial conditions. A further remarkable feature is that, in the unpredictable phase (many agents), adding randomness in the way agents chose, collective fluctuations decrease. This is at odd with conventional wisdom according to which collective fluctuations increase when stochasticity at the micro scale – e.g. 'temperature' – increases. Even more remarkable is the fact that in the predictable phase (few agents) collective fluctuations do not depend at all from microscopic randomness.

All these features find their explanation together with a precise quantitative assessment within a stochastic theory of the Minority Game [Marsili and Challet, 2001a]. Microscopic 'temperature' affects learning rates and finally turns out to act as the inverse of a global temperature, thus explaining the unconventional dependence of fluctuations on microscopic noise. This approach, based on the neglect of time dependent volatility fluctuations, also unveils the peculiar interconnection between initial conditions and correlated fluctuations, leaving us with a complete picture of the Minority Game behavior. Although an approximation, this theory turns out to be quite precise in general, except close to the phase transition in finite sytems.

We shall finally comment on a powerful alternative approach, based on the generating functional [Heimel and Coolen, 2001]. The extraction of quantitative results from this approach is made difficult by the mathematical complexity of the resulting equations. On the other hand, this method is exact. Even if it does not yet provide a complete picture of how Minority Game works, the generating function is a very promising approach.

[1] I.e. the frequencies with which agents play their strategies.

2 Minority and Majority Games as Market models

The minority game has been proposed[Challet and Zhang, 1997, Zhang, 1998] to model speculative behavior in financial markets. Agents sell and buy asset shares or currencies with the only goal of profiting from market's fluctuations. The minority game is a highly simplified picture of this context where agents can take, in each period, one of two actions. The agents who take the minority action win, whereas the majority looses. The connection with financial markets is established naively by observing that markets are instruments for reallocating goods. No gain can be made, in principle, by pure trading. Hence the market should be a zero sum game. Transaction costs and other frictions make it an unfavorable game, on average, i.e. a minority game[2].

A naïve argument, however, suggests the opposite conclusion: when everybody is going to buy the price will raise and hence it is convenient to buy. Hence speculative markets should be rather similar to majority games. One may argue that only the minority of agents who buy first win whereas the other lose and eventually enter into endless arguments.

The problems with arguments in support of either the minority or the majority game essentially arise from the difficulties related with the definition of the payoff of a single transaction. Strictly speaking, buying or selling an asset does not change agent's wealth but just the composition of his portfolio. A real gain can be measured only if after buying or selling something at time t, the same is sold or bought again at time $t' > t$. This leaves the trader with the same amount of asset but with a gain or a loss of money. Therefore market gains depends on more than one time and in general, assessing the performance of a trading strategy is a complex inter-temporal problem. The ultimate conclusion of these arguments is that financial markets cannot be described by simple markovian agent based models, but rather require rather sophisticated intertemporal models of traders' behavior.

It may be reasonable to assume that traders first form an expectation on the behavior of the market and then optimize their behavior with respect to this expectation. Expectations are eventually revised and modified, on a longer time-scale, if they contrast with actual market behavior.

We show indeed that the minority game can be derived assuming that agents follow this behavior, from a market mechanism. From this viewpoint, expectations lie at the very basis of the definition of the minority game as a market model. Depending on the expectations of agents we can distinguish between *fundamentalists* or *contrarian* traders – who perceive the market as a minority game – and *trend followers* – who perceive it as a majority game.

Before entering into details, let us mention that Ref. [Marsili, 2001] analyzes a simple case where both fundamentalists and trend followers are present in the market. Since the two groups have opposite expectations on the price process, the relevant question is which of these expectations is fulfilled by the actual price process? It turns out that the price process satisfies the expectations of whatever group is more numerous: In a market with a majority of fundamentalists the minority mechanism will prevail and the expectations of fundamentalists will be satisfied. On the contrary, if trend followers are the majority their expectations

[2]Ref. [Challet et al., 2000a] expands these types of arguments in much more details.

will be satisfied and price process will acquire a trend. This agrees with the well known observation that expectations of traders in a market can be self-fulfilling prophecies.

Let us imagine that time t is discrete, there are N agents and they submit all together their orders $a_i(t)$ to the market ($i = 1, \ldots, N$).

The single time step is split in three phases:

<u>time $t-\epsilon$</u> Agents take their choices on the basis of their accumulated experience up to time $t-1$ and submit their orders $a_i(t)$.

<u>time t</u> The market aggregates orders $a_i(t)$ from traders and forms a price $p(t)$.

<u>time $t+\epsilon$</u> Agents learn: they update their experience by evaluating the success of their actions. This will enter into the decision process at time $t+1-\epsilon$.

Agents do not know the price at which the transaction will actually take place. Secondly, given that agents cannot define a real payoff for the present transaction as discussed above, they have to resort to "perceived" or "expected" payoffs in the learning phase.

We shall discuss later how agents take their decisions and how they update their behavior on the basis of perceived payoffs. For the moment being let us focus on the second step and define the market interaction in detail. Let $a_i(t) > 0$ mean that agent i contributes $a_i(t)$\$ to the demand for the asset. Likewise $a_i(t) < 0$ means that i sells $-a_i(t)/p(t-1)$ units of asset, which is the current equivalent of $|a_i(t)|$\$. With $a_i(t) = \pm 1$ and $A(t) = \sum_i a_i(t)$, the demand is given by $D(t) = \frac{N+A(t)}{2}$, whereas the supply is $S(t) = \frac{N-A(t)}{2p(t-1)}$. Price is fixed by the market clearing condition, $p(t) = D(t)/S(t)$, i.e.

$$p(t) = p(t-1)\frac{N+A(t)}{N-A(t)}. \tag{1}$$

Consider an agent who buys 1\$ of asset at time t (i.e. $a_i(t) = 1$): He exchanges 1\$ with $1/p(t)$ units of asset. Was the choice $a_i(t) = 1$ the "best" one?

In order to answer this question, we may imagine that agent i considers selling $1/p(t)$ units of assets at time $t+1$. This leads to a payoff

$$u_i(t) = \frac{p(t+1)}{p(t)} - 1, \qquad \text{if } a_i(t) = +1 \tag{2}$$

However, the price $p(t+1)$ will only be revealed after agents communicate their investments decisions $a_j(t+1)$ for all j. If agents want to use Eq. (2) to revise their choice rule before deciding $a_i(t+1)$, they have to replace $p(t+1)$ in it by their expectation at time t, denoted by $E_t^{(i)}[p(t+1)]$. Let us assume that:

$$E_t^{(i)}[p(t+1)] = (1-\psi_i)p(t) + \psi_i p(t-1) \tag{3}$$

Then $E_t^{(i)}[u_i(t)|a_i(t) = +1] = -\psi_i[p(t) - p(t-1)]/p(t)$ and, using Eq. (1) we find $E_t^{(i)}[u_i(t)|a_i(t) = +1] = -2\psi_i A(t)/[N + A(t)]$.

Likewise, if agent i sells $1/p(t-1)$ units of assets at time t (i.e. $a_i(t) = -1$) and buys it back at the expected price $E_t^{(i)}[p(t+1)]$, elementary algebra leads to $E_t^{(i)}[u_i(t)|a_i(t) = -1] = 2\psi_i A(t)/[N - A(t)]$. This means that:

$$u_i[a_i(t), A(t)] = E_t^{(i)}[u_i(t)] = -2\psi_i a_i(t)\frac{A(t)}{N + a_i(t)A(t)}. \quad (4)$$

Notice that agents who took the majority action $a_i(t) = \text{sign}\, A(t)$ "receive a payoff" $-2\psi_i|A(t)|/[N + |A(t)|]$ whereas agents in the minority get $2\psi_i|A(t)|/[N - |A(t)|]$. If $\psi_i > 0$ the minority is the winning side and indeed Eq. (4) reduces to the usual payoffs of the minority games [3]. Agents with $\psi_i > 0$ may be called *fundamentalists* as they believe that market prices fluctuate around a fixed value, so that future price is an average of past prices. They may also be called *contrarians* since they believe that the future price increment $\Delta p(t+1) = p(t+1) - p(t)$ is negatively correlated with the last one

$$E_t^{(i)}[\Delta p(t+1)] = -\psi_i \Delta p(t).$$

On the other hand, if $\psi_i < 0$ the game turns into a majority game. More precisely, agent i perceives the game as one in which he prefers to stay in the majority. These type of agents may be called trend followers since they believe that future price increments $\Delta p(t+1)$ are positively correlated with past ones, as if the price were following a monotonic trend.

3 Definition of the Minority Game

We focus on the simplest non-trivial version of the model to make the derivation as straight and simple as possible. The derivation can be extended in a straightforward manner to a number of variations on the theme, adding frills here and there on the skeleton Minority Game which we discuss. These extensions are discussed in the final section. The overall picture turns out to be quite robust.

3.1 From the MG definition to a computer code

The MG was introduced in Ref. [Challet and Zhang, 1997] with the following definition:

> Let us consider a population of N (odd) players, each has some finite number of strategies S. At each time step, everybody has to choose to be in side A or side B. The payoff of the game is to declare that after everybody has chosen side independently, those who are in the minority side win. In the simpliest version, all winners collect a point. The players make decisions based on the common knowledge of the past record. We further limit the record to contain only yes and no e.g. the side A is the winning side or not, without the actual attendance number. Thus the system's signal can be represented by a binary sequence, meaning A is the winning side (1) or not (0).
>
> Let us assume that our players are quite limited in their analysing power, they can only retain last M bits of the system's signal and make their next decision basing only on these M bits. Each player

[3] It turns out that, in the minority game ($\psi_i > 0\ \forall i$) $A(t) \sim \sqrt{N}$ is negligible compared to N in the denominator and it can be dropped. Then one recovers the linear payoffs used in Refs. [Challet and Marsili, 1999, Cavagna et al., 1999, Challet et al., 2000c, Marsili et al., 2000].

has a finite set of strategies. A strategy is defined to be the next action (to be in A or B) given a specific signal's M bits. The example of one strategy is illustrated in table 1 for M=3.

signal	prediction
000	1
001	0
010	0
011	1
100	1
101	0
110	1
111	0

There are 8 ($= 2^M$) bits we can assign to the right side, each configuration corresponds a distinct strategy, this makes the total number of strategy to be $2^{2^M} = 256$. This is indeed a fast increasing number, for $M = 2, 3, 4, 5$ it is 16, 256, 65536, 65536^2. We randomly draw S strategies for each player, and some strategies maybe by chance shared. However for moderately large M, the chance of repetition of a single strategy is exceedingly small. Another special case is to have all 1's (or 0's) on the RHS of the table, corresponding to the fixed strategy of staying at one side no matter what happens.

Let us analyse the structure of this Minority game to see what to expect. Consider the extreme case where only one player takes a side, all the others take the other. The lucky player gets a reward point, nothing for the others. Equally extreme example is that when $(N-1)/2$ players in one side, $(N+1)/2$ on the other. From the society point of view, the second situation is preferable since the whole population gets $(N-1)/2$ points whereas in the first example only one point—a huge waste. Perfect coordination and timing would approach the 2nd, disaster would be the first example. In general we expect the population to behave between the above two extremes.

This binary game can be easily simulated for a large population of players. Initially, each player draws randomly one out of his S strategies and use it to predict next step, an artificial signal of M bits is also given. All the S strategies in a player's bag can collect points depending if they would win or not given the M past bits, and the actual outcome of the next play. However, these points are only *virtual* points as they record the merit of a strategy as if it were used each time. The player uses the strategy having the highest accumulated points (capital) for his action, he gets a real point only if the strategy used happens to win in the next play.

As a first step to understand the MG, let us try to translate this wordy definition into a computer program. The result may look like the following[4]:

[4] We use FORTRAN language with operators `.eq.` and `.gt.` replaced by = and > respectively for the sake of readability.

```
    ...
C AGENTS' CHOICE
    do i=1,N
      do s=1,S
        if(points(i,s) > points(i,bestStrategy(i))) bestStrategy(i) = s
      end do
    end do
C MARKET INTERACTION
    N1=0
    do i=1,N
      N1 = N1 + side(i,bestStrategy(i),mu)
    end do
    if (N1 > N/2) then
      winSide = 0
    else
      winSide = 1
    end if
C LEARNING
    do i=1,N
      do s=1,S
        if (side(i,s,mu) = winSide) points(i,s) = points(i,s)+1
      end do
    end do
C INFORMATION UPDATE
    mu = mod(2*mu + winSide, 2**M)
    ...
```

This code is a good starting point to analyze the dynamics: First agents choose the strategy they play – stored in the variable `bestStrategy(i)` – looking at the virtual points `points(i,s)` that each strategy `s` has accumulated. Agents pick the strategy with the largest number of points[5]. The side prescribed by this strategy, given the sequence `mu` of the last M outcomes, is stored in the table `side(i,s,mu)` for each agent `i` and each strategy `s`. Note that this depends on `mu` which encodes the history of recent games. The tables `side(·,·,·)` are drawn at random – with `side(i,s,mu)=0` or `1` with equal probability – at the beginning of the game, i.e. in the initialization section of the program. Once every agent has fixed his `bestStrategy(i)`, the attendance of the two sides is computed: `N1` is the number of agents who took the choice `1` and `winSide` is the winning side: `winSide=1` if `N1<N/2` and `winSide=0` otherwise (we assume `N` is odd). Given the outcome `winSide` of the game, agents updated their scores `points(i,s)`. Finally the history `mu` – which is an integer variable – is updated for the next time-step by the last instruction. In this way the first `M` bits in the binary representation of `mu` are the last `M` values of `winSide`.

The central quantity of interest is the difference in the attendance of the two sides:

$$A = 2 * \mathtt{N1} - \mathtt{N} \tag{5}$$

But before discussing the fluctuations of A, let us review the key steps which have led to an analytically tractable version of the MG.

[5] In case of ties a tie-breaking rule has to be decided. The one used here is an example. Other rules can be used without affecting much the collective properties.

3.2 Simplifying the Minority Game dynamics

In its simplicity the Minority Game as defined above, captures quite complex phenomena. The route to a thorough analytical approach have been made possible by simplifying the model still further, while preserving its rich dynamical behaviour.

1. The first has been the observation by Cavagna [Cavagna, 1999] that the fluctuations of A are left largely unaffected if the dynamics of the history

$$\texttt{mu} = \texttt{mod}(2 * \texttt{mu} + \texttt{winSide}, 2 * *\texttt{M}) \tag{6}$$

is replaced by a random draw from the integers $\{0, \ldots, 2^M - 1\}$ with uniform probability:

$$\texttt{mu} = \texttt{int}(2 * *\texttt{M}\ \texttt{rand}()). \tag{7}$$

While μ in Eq. (6) encodes the *real* history, in Eq. (7) μ is just a random piece of information[6].

One of the ideas behind Eq. (6) was to describe a closed system where agents process and react to a piece of information they themselves produce. The results of Ref. [Cavagna, 1999] suggests that this feedback is largely irrelevant. The *endogenous* information process of Eq. (6) may well be replaced by Eq. (7) which models a news arrival process of *exogenous* information. With exogenous information, the number of values that μ can take is not restricted to be a power of two. We call this number P henceforth, and we shall have

$$P \equiv 2^M$$

for endogenous information.

A close inspection of Fig 1 of Ref. [Cavagna, 1999] shows that the conclusion on the irrelevance of the origin of information does not hold exactly [Challet and Marsili, 2000]. We shall go back to this issue in the last section 7. At any rate, the passage from endogenous (Eq. 6) to exogenous (Eq. 7) information represents a great simplification of the model.

2. A further simplification of the original model is to replace the accounting of the points `points(i,s)` by a linear dynamics of scores $U_{i,s}(t)$. It was noticed in Refs. [Johnson et al., 1998, Challet and Marsili, 1999, Cavagna et al., 1999] and later shown in Ref. [Challet et al., 2000d], that this modification does not alter the qualitative behaviour of the model. In practice, this amounts to replacing the update of points `points(i,s)` by

$$U_{i,s}(t+1) = U_{i,s}(t) - a_{i,s}^{\mu(t)} \frac{A(t)}{N} \tag{8}$$

where we introduced the convenient notation

[6] Early works (e.g. [Challet and Zhang, 1997, Cavagna, 1999]) refer to *memory* rather than to *history*. We prefer the latter term because, strictly speaking, the memory of agents is stored in `points` rather than in μ.

$$a_{i,s}^{\mu} = 2*\texttt{side(i,s,}\mu) - 1 \tag{9}$$

for the strategy tables. Again the strategy s is rewarded [i.e. $U_{i,s}(t+1) - U_{i,s}(t) > 0$] when it predicts correctly the minority sign, i.e. if $a_{i,s}^{\mu} = -\text{sign}\, A(t)$, and penalised otherwise.

3. Finally it is convenient to generalize in a probabilistic fashion the way in which agents take decisions. In the original Minority Game the strategy $s_i(t)$ which agent i uses at time t is that with the highest score:

$$s_i(t) = \arg\max_s U_{i,s}(t). \tag{10}$$

This introduces a mathematical discontinuity which is hard to deal with analytically. Ref. [Cavagna et al., 1999] suggested to overcome this difficulty by resorting to a probabilistic choice model:

$$\text{Prob}\{s_i(t) = s\} = \frac{e^{\Gamma U_{i,s}(t)}}{\sum_{s'} e^{\Gamma U_{i,s'}(t)}} \tag{11}$$

with $\Gamma > 0$. This is reminiscent of the Gibbs distribution for physicists [Yeomans, 1992] and Γ appears as an "individual inverse temperature" – whereby the name of Thermal Minority Game [Cavagna et al., 1999]. Eq. (11) is also a very well known choice model among economists, known as the *Logit model*[7] [Luce, 1959, McFadden, 1981].

[7] The probabilistic nature of agents' choice does not necessarily imply that the agents randomise their behaviour on purpose. McFadden has indeed shown that Eq. (11) models individuals who maximise an "utility" which has an implicit random idiosyncratic part $\eta_{i,s}$:

$$s_i(t) = \arg\max_s [\Gamma U_{i,s}(t) + \eta_{i,s}(t)] \tag{12}$$

The constant Γ is the relative weight which agents assign to the empirical evidence accumulated in $U_{i,s}$ with respect to random idiosyncratic shocks $\eta_{i,s}$. If $\Gamma \to \infty$ agents always play their best strategy according to the scores, while if Γ decreases agents take less into account past performances. For a generic distribution $P_\eta(x) \equiv \text{Prob}\{\eta_{i,s}(t) < x\}$ we have

$$\text{Prob}\{s_i(t) = s\} = \int_{-\infty}^{\infty} dP_\eta(x) \prod_{r \neq s} P_\eta[x + \Gamma(U_{i,s} - U_{r,i})].$$

This coincides with Eq. (11) if $P_\eta(x) = \exp(-e^{-x})$ is the Gumbel distribution [McFadden, 1981]. Different distributions of $\eta_{i,s}$ lead to choice models which are different from Eq. (11), but the model of Eq. (11) is unique in that it satisfies the axiom of *independence from irrelevant alternatives*: this states that the relative odds of choices s and s' does not depend on whether another choice s'' is possible or not. In addition, there is a natural derivation of the Gumbel distribution in the case where agents want to maximize an utility function $W_i(s, \sigma_1, \ldots, \sigma_n) = U_{i,s} + nV_i(\sigma_1 \ldots, \sigma_n|s)$ which depends also on n variables σ_k which take g values each. The factor n in the second term implies that the weight of all variables is equivalent in decision making. The choice behavior with respect to the other variables, which is not of our explicit interest, is modelled in a probabilistic way. With an opportune choice of $U_{i,s}$, let us assume that $V_i(\sigma_1 \ldots, \sigma_n|s)$ are i.i.d. gaussian variables with zero mean and variance v. Then, using extreme value statistics [Galambos, 1987]

$$\eta_{i,s} = \max_{\{\sigma_k\}} nV_i(\sigma_1 \ldots, \sigma_n|s) = V_0 + \sqrt{\frac{n}{2v \log g}} Y_{i,s}$$

where V_0 is a uninfluential constant and $Y_{i,s}$ is distributed according to the Gumbel distribution. This suggests that $\Gamma = \sqrt{(2v \log g)/n}$, i.e. when there are many other choices ($n \gg 1$)

Summarizing, the dynamics of the simplified Minority Game is described by the following equations:

$$\begin{aligned}
\text{Prob}\{s_i(t)=s\} &= \frac{e^{\Gamma U_{i,s}(t)}}{\sum_{s'} e^{\Gamma U_{i,s'}(t)}} && \text{(choices of agents)} && (13)\\
\text{Prob}\{\mu(t)=\nu\} &= \tfrac{1}{P}, \quad \nu=1,\ldots,P && \text{(choice of Nature)} && (14)\\
A(t) &= \textstyle\sum_{i=1}^{N} a_{i,s_i(t)}^{\mu(t)} && \text{(market aggregation)} && (15)\\
U_{i,s}(t+1) &= U_{i,s}(t) - a_{i,s}^{\mu(t)}\tfrac{A(t)}{N} && \text{(learning)} && (16)
\end{aligned}$$

The strategies $a_{i,s}^{\mu}$ are randomly drawn at the beginning of the game and then they are kept fixed. Hence they can be considered as fixed (quenched) disorder.

There is a further simplification which does not entail a modification of the dynamics but just a restriction. Early numerical studies have shown that varying the number S of strategies given to each agent, the MG's behaviour remains qualitatively the same. Actually cooperative effects manifest most strongly for $S=2$ strategies.

Hence it is preferable to restrict attention to the case where all agents have $S=2$ strategies (the case with $S>2$ strategies will be discussed in section 7). Ref. [Challet and Marsili, 1999] introduced a convenient notation by labelling strategy 1 as -1 and strategy 2 as $+1$ so that agent i controls the variable s_i which takes values ± 1; such a variable is called a spin in physics. The strategies of agent i can be decomposed

$$a_{i,s}^{\mu} = \omega_i^{\mu} + s_i \xi_i^{\mu}$$

into a constant $\omega_i^{\mu} = \frac{a_{+,i}^{\mu}+a_{-,i}^{\mu}}{2}$ and a variable component, with $\xi_i^{\mu} = \frac{a_{+,i}^{\mu}-a_{-,i}^{\mu}}{2}$. In binary games $|a_i^{\mu}|=1$, therefore ω_i^{μ}, $\xi_i^{\mu} = 0, \pm 1$ but $\omega_i^{\mu}\xi_i^{\mu} = 0$. The fact that for some μ, agent i have strategies which prescribe the same action ($\xi_i^{\mu}=0$ and $\omega_i^{\mu} \neq 0$) lies at the origin of cooperation in the game (see section 7). This decomposition allows us to express $A(t)$ in a form where its dependence on the quenched disorder variables $\omega_i^{\mu}, \xi_i^{\mu}$ and on the dynamical variables $s_i(t)$ is made explicit. Indeed Eq. (15) becomes

$$A(t) = \Omega^{\mu(t)} + \sum_{i=1}^{N} \xi_i^{\mu(t)} s_i(t), \quad \text{with} \quad \Omega^{\mu} = \sum_{i=1}^{N} \omega_i^{\mu}. \tag{17}$$

A further simplification arises from the fact that, only the difference between the scores of the two strategies is important in the dynamics. In other words, each agent can be described in terms of a single dynamical variable

$$Y_i(t) = \Gamma \frac{U_{+,i}(t) - U_{-,i}(t)}{2}. \tag{18}$$

The probability distribution of $s_i(t)$ becomes

$$\text{Prob}\{s_i(t) = \pm 1\} = \frac{1 \pm \tanh Y_i(t)}{2}. \tag{19}$$

σ_k which are affected by the choice of s we expect a small value of Γ and *viceversa*.

Finally, the dynamics of $Y_i(t)$ is derived taking the difference of Eqs. (16) for $s = \pm 1$:

$$Y_i(t+1) = Y_i(t) - \frac{\Gamma}{N} \xi_i^{\mu(t)} A(t). \tag{20}$$

The model defined in Eqs. (14,17,19) and (20) shall be our reference model for the remaining of this chapter. Below we shall discuss its generic behaviour and the theoretical approach based on statistical physics which has allowed to understand it.

3.3 Some convenient notations

Before entering into the details, it is convenient to discuss statistical averages and to introduce the relative notations. There are two sources of randomness in the MG. One is the choice of information $\mu(t)$ and the other is agents' choice of strategies $s_i(t)$. We shall be interested in the stationary state of the game[8]. Hence, for any quantity $Q(t)$, we denote by

$$\langle Q \rangle = \lim_{t_0 \to \infty} \lim_{T \to \infty} \frac{1}{T} \sum_{t=t_0+1}^{t_0+T} Q(t) \tag{21}$$

its average value in the stationary state. We shall also be interested in conditional averages for a particular value of the information μ. We denote it as

$$\langle Q|\mu \rangle = \lim_{t_0 \to \infty} \lim_{T \to \infty} \frac{1}{T_\mu} \sum_{t=t_0+1}^{t_0+T} Q(t) \delta_{\mu(t),\mu} \quad \text{where } T_\mu = \sum_{t=t_0+1}^{t_0+T} \delta_{\mu(t),\mu}. \tag{22}$$

Finally, averages over μ will be denoted by an over-line:

$$\overline{Q} = \frac{1}{P} \sum_{\mu=1}^{P} Q^\mu$$

This notation shortens considerably some expressions. We implicitly assume that averages under the over-line are conditional on μ. So, for example

$$\overline{F(\langle Q \rangle)} = \frac{1}{P} \sum_{\mu=1}^{P} F(\langle Q|\mu \rangle).$$

Note that clearly $\langle Q \rangle = \overline{\langle Q \rangle}$ but $\langle Q \rangle^2 \neq \overline{\langle Q \rangle^2}$.

4 Generic behaviour of the MG

Let us rephrase in mathematical terms the generic behaviour of the MG sketched in the previous chapter. Early papers focused on the cooperative properties of the system in the stationary state. Symmetry arguments suggest that none of

[8] That a stationary state exists for the Minority Game exists can be shown observing that *i)* it is a Markov process *ii)* it can be approximated with a finite space Markov process (discretizing $Y_i(t)$ on a grid of Λ points) to an arbitrary precision (letting $\Lambda \to \infty$).

the two groups 0 or 1 will be systematically the minority one. This means that $A(t)$ will fluctuate around zero, i.e. $\langle A\rangle = 0$. The size of fluctuations of $A(t)$, instead, displays a remarkable non-trivial behaviour. The variance

$$\boxed{\sigma^2 \equiv \langle A^2\rangle} \tag{23}$$

of $A(t)$ in the stationary state is – quoting from Ref. [Savit et al., 1999] – "a convenient reciprocal measure of how effective the system is at distributing resources". The smaller σ^2 is, the larger a typical minority group is. In other words σ^2 is a reciprocal measure of the *global efficiency* of the system. This is obvious if the payoff function is linear, as in Eqs (16) and (8): in that case the total payoff given to the agents $-\sum_i a_i(t)A(t) = -A^2(t)$, hence, σ^2 measures the average total loss of the agents per time-step.

Early numerical studies [Challet and Zhang, 1997, Savit et al., 1999, Challet and Zhang, 1998] uncovered a remarkably rich phenomenology as a function of M and the number of agents N. Savit *et al.* [Savit et al., 1999] found that the collective behaviour does not depend independently on M and N but only on the ratio

$$\boxed{\alpha \equiv \frac{2^M}{N} = \frac{P}{N}} \tag{24}$$

between the number $P = 2^M$ of possible histories μ and the number of agents, as illustrated in Fig. 1. This means that typically $A(t) \propto \sqrt{N}$ for fixed α. When $\alpha \gg 1$ information is too complex and agents essentially behave randomly. Indeed σ^2/N converges to one – the value it would take if agents were choosing their side by coin tossing. As α decreases – which means that M decreases or the number of agents increases – σ^2/N decreases suggesting that agents manage to exploit the information in order to coordinate. But when agents become too numerous, σ^2/N starts increasing with N. Savit *et al.* [Savit et al., 1999] found that, at M fixed, σ^2 increases linearly with N as long as $N \ll P$ but with a quadratic law $\sigma^2 \sim N^2$ for $N \gg P$, which implies $\sigma^2 \sim 1/\alpha$ for $\alpha \ll 1$. The behaviour for $\alpha \ll 1$ has been attributed to the occurrence of 'crowd effects', and it has been studied in some detail both numerically and by approximate methods [Hart et al., 2001, Zhang, 1998, Challet and Zhang, 1998, D'hulst and Rodgers, 1999, Manuca et al., 2000, Heimel and Coolen, 2001, Caridi and Ceva, 2003].

A further interesting observation of Savit and co-workers [Savit et al., 1999] comes from their analysis of the probability

$$p^\mu = \langle \theta(A)|\mu\rangle$$

that the minority is on one given side conditional on the value of μ (here $\theta(x) = 0$ for $x < 0$ and 1 otherwise). This is an important quantity as it tells us whether $\mu(t)$ carries some information on the attendance $A(t)$ or not. Savit *et al.* observed that $p^\mu = 1/2$ for $\alpha \ll 1$: the minority was falling on either side with equal probability irrespective of μ. But when $\alpha \gg 1$ the minority happens to be more likely on one side (i.e. $p^\mu \neq 1/2$), depending on the value of μ. This means that the value of $\mu(t)$ contains some information on $A(t)$ because it makes possible a better than random prediction of the minority side.

These observations have been sharpened in Ref. [Challet and Marsili, 1999] by confirming the existence of a phase transition located at the point where σ^2 attains its minimum ($\alpha_c \approx 0.34$ for $S = 2$). The transition separates a

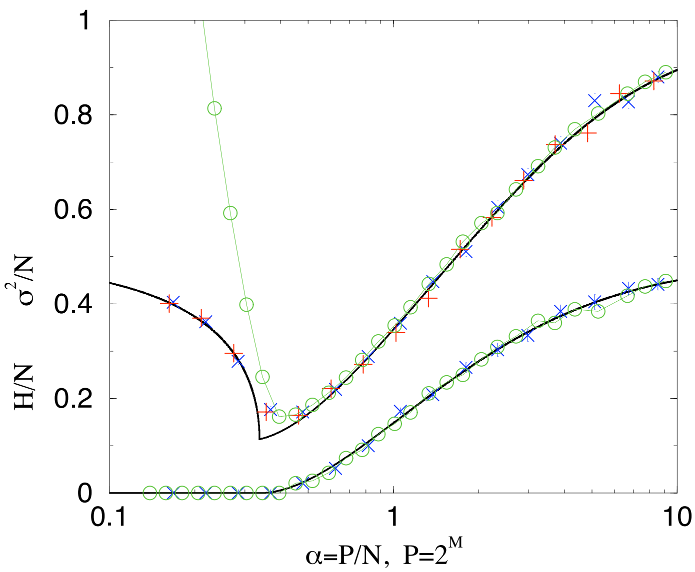

Figure 1: Global efficiency σ^2/N for $P = 128$, and 256 and N ranging from $0.1P$ to $10P$. Data set with $\Gamma = 0.1$ ($+$ $P = 128$ and $\times$ $P = 256$) or $\Gamma = \infty$ are shown ($\circ$ $P = 128$). Agents have two trading strategies. Each point is obtained as the average over 100 random systems in the asymptotic stationary state of the dynamics. For large N the collective properties only depend on the ratio $\alpha = P/N$. For small values of α, σ^2 is larger for fast learning ($\circ$ correspond to $\Gamma = \infty$) than for slow learning ($\times$, $+$ correspond to $\Gamma = 0.1$). The full line are the results of the theory based on the statistical mechanics approach, which is valid in the limit $P \to \infty$ and for small Γ. The predictability H/N as a function of α, for the same systems as above, is also shown at the bottom of the figure. H does not depend on Γ.

symmetric ($\alpha < \alpha_c$) from an asymmetric phase ($\alpha > \alpha_c$). The symmetry which is broken is that of the average $\langle A|\mu\rangle$ of $A(t)$ conditional on the history μ.
In the asymmetric phase, $\langle A|\mu\rangle \neq 0$ for at least one μ. Hence knowing the history $\mu(t)$ at time t, makes the sign of $A(t)$ statistically predictable. A measure of the degree of predictability is given by the function

$$\boxed{H = \frac{1}{P}\sum_{\mu=1}^{P}\langle A|\mu\rangle^2 = \overline{\langle A\rangle^2}.} \tag{25}$$

In the symmetric phase $\langle A|\mu\rangle = 0$ for all μ and hence $H = 0$.

Ref. [Challet and Marsili, 1999] found that, at fixed M or P, H is a decreasing function of the number N of agents[9] (see Fig. 1): This means that newcomers exploit the predictability of $A(t)$ and hence reduce it. The same article also introduced the concept of *frozen agents*, who always play the same

[9] A slightly different quantity $\theta = \frac{1}{P}\sum_{\mu=1}^{P}\langle \text{sign}\,(A)|\mu\rangle^2 = \overline{(2p-1)^2}$ was actually studied in [Challet and Marsili, 1999].

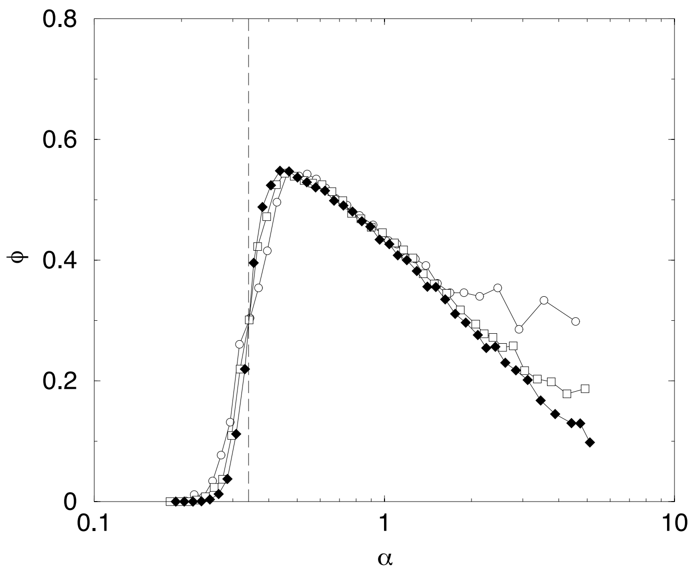

Figure 2: Fraction of frozen agents versus $\alpha = P/N$ for $M = 6$ (circles), 7 (squares) and 8 (diamonds). The critical point is located at the intersection of the three curves.

strategy. The fraction ϕ of frozen agents is reported in Fig. 4 and it has a discontinuity which provides a very precise determination of α_c.

Finally Ref. [Cavagna et al., 1999] observed that fluctuations σ^2 increase with Γ. That is a quite remarkable finding specially if one considers Γ as an inverse temperature as suggested by Eq. (13): Usually fluctuations decrease when the temperature decreases, whereas σ^2 increases when the 'temperature' $1/\Gamma$ decreases. Note that Ref. [Cavagna et al., 1999] also reports a rise in σ^2 for very high 'temperatures' (i.e. for $\Gamma \ll 1$). This was later found – first in Ref. [Bottazzi et al., 2003] then in Ref. [Challet et al., 2000b] – to be due to lack of equilibration in numerical simulations.

What is also quite remarkable is that, as shown in Fig. 1, σ^2 depends on Γ only for $\alpha < \alpha_c$ but not for $\alpha > \alpha_c$. Nor does H for all values of α. Furthermore, the stationary state depends on initial conditions for $\alpha < \alpha_c$ [Challet et al., 2000c, Garrahan et al., 2000, Marsili and Challet, 2001b, Marsili and Challet, 2001a, Marsili, 2001, Heimel and Coolen, 2001]: the larger the spread or the asymmetry in the initial conditions $\{Y_i(0)\}$ the smallest the value of σ^2. This dependence disappears for $\alpha > \alpha_c$ and the dynamics 'forgets' about initial conditions, as in ergodic systems (see subsection 6.4).

These results leave us with a number of open questions that a theory of the Minority Game should address.

5 Minority Game without information

Interestingly, a hint on some of these answers can be obtained from the study of an highly simplified version of the MG [Marsili and Challet, 2001b, Marsili, 2001] where the collective behavior can be easily understood with simple mathematics.

In the minority game the quantities of interest are the first two moments of $A(t)$:

$$\langle A\rangle = \lim_{t_0,T\to\infty}\frac{1}{T}\sum_{t=t_0+1}^{t_0+T}A(t)$$

$$\sigma^2 \equiv \langle A^2\rangle = \lim_{t_0,T\to\infty}\frac{1}{T}\sum_{t=t_0+1}^{t_0+T}A^2(t).$$

If agents are rational at all, we expect that they will drive the system to a state where none of the two actions $a_i = \pm 1$ identifies systematically the minority side. Hence, we expect $\langle A\rangle = 0$. σ^2 gives instead a measure of the efficiency of the systems because it tells how many more losers than winners are there. For illustrative purposes, let us compare the state where $A(t) = 0\ \forall t$ to the state where $A(t) = (-1)^t N$. Both have $\langle A\rangle = 0$ however in the former no agent loses ($\sigma^2 = 0$) whereas in the latter all agents lose ($\sigma^2 = N^2$).

Agents learn from past experience which action $a_i(t)$ is the best one. The learning dynamics is the one used in general in minority games and it is well rooted in the economic literature [Rustichini, 1999]. The past experience of agent i is stored in the "score" $\Delta_i(t)$: $\Delta_i(t) > 0$ means that the action $a_i = +1$ is (perceived as) more successful than $a_i = -1$ and *vice-versa*. Agents use the information accumulated in $\Delta_i(t)$ to take decisions[10]:

$$\mathrm{Prob}\{a_i(t) = \pm 1\} \equiv \frac{e^{\Delta_i(t)}}{e^{\Delta_i(t)} + e^{-\Delta_i(t)}} \tag{26}$$

and they update $\Delta_i(t)$ by

$$\Delta_i(t+1) = \Delta_i(t) - \Gamma\frac{A(t)}{N} \tag{27}$$

This learning dynamics is easily understood: if $A(t) < 0$ agents observe that the best action was $+1$ at time t. Hence they increase Δ_i and the probability of playing $a_i = +1$ (see Eq. (26)). The parameter Γ modulates the strength of the response in the behavior of agents to the "stimulus" $A(t)/N$. Let us finally assume that the initial conditions $\Delta_i(0)$ are drawn from a distribution $p_0(\Delta)$ with a standard deviation s. How does the collective behavior depends on the parameters Γ and s?

Notice that $y(t) = \Delta_i(t) - \Delta_i(0)$ does not depend on i, for all times. For $N \gg 1$, the law of large numbers allows us to approximate $A(t)/N$ by its average value[11]. In Eq. (27), this yields a dynamical equation for $y(t)$:

$$y(t+1) \cong y(t) - \Gamma\langle\tanh[y(t) + \Delta(0)]\rangle_0 \tag{28}$$

[10]The exponential form, which results from a Logit discrete choice model, is taken here for simplicity. Any increasing continuous function $\chi_i(x)$, with $0 \le \chi_i(x) \le 1$ for all real x, $\chi(x) \to 0$ as $x \to -\infty$ and $\chi(x) \to 1$ as $x \to \infty$, leads to the same results [Marsili and Challet, 2001b].

[11]The average value of $A(t)/N$ is computed using Eq. (26), which gives $\langle a_i(t)\rangle = \tanh \Delta_i(t)$, and averaging over the distribution $p_0(\Delta)$ of initial conditions $\Delta_i(0)$.

where the average $\langle \ldots \rangle_0$ is on the distribution p_0 of initial conditions. This equation admits a fixed point solution $y(t) = y^*$ for all t, where y^* is the solution of

$$\langle \tanh[y^* + \Delta(0)] \rangle_0 = \langle A \rangle = 0 \tag{29}$$

If this solution is stable, the distribution of relative scores $\Delta_i(t)$ shift bodily from the initial conditions and settles around the origin, in order to satisfy $\langle A \rangle = 0$ (Eq. (29)). Given that $\langle a_i \rangle = \tanh[y^* + \Delta_i(0)]$, it is not difficult to find that

$$\sigma^2 = \sum_{i=1}^{N} (1 - \langle a_i \rangle^2) = N \left\{ 1 - \langle \tanh[y^* + \Delta(0)]^2 \rangle_0 \right\}$$

for this state. Notice that $\sigma^2 \propto N$ and it *decreases* with the spread of the distribution of initial conditions.

When is this a stationary state of the dynamics? To answer this question it suffices to study the linear stability of the dynamics. We set $y(t) = y^* + \delta y(t)$ and expand Eq. (28) to linear order. It is easy to find that the fixed point y^* is stable only for

$$\Gamma < \Gamma_c = \frac{2}{1 - \langle \tanh[y^* + \Delta(0)]^2 \rangle_0} = \frac{2N}{\sigma^2}. \tag{30}$$

When $\Gamma > \Gamma_c$ we find periodic solutions $y(t) = y^* + z^*(-1)^t$ where y^* and z^* satisfy

$$\frac{\langle \tanh[y^* + z^* + \Delta(0)] \rangle_0 + \langle \tanh[y^* - z^* + \Delta(0)] \rangle_0}{2} = 0 \tag{31}$$

$$\frac{\langle \tanh[y^* + z^* + \Delta(0)] \rangle_0 - \langle \tanh[y^* - z^* + \Delta(0)] \rangle_0}{2} = \frac{2z}{\Gamma}. \tag{32}$$

The parameter z^* plays the role of an order parameter of the transition at Γ_c ($z^* = 0$ for $\Gamma < \Gamma_c$). Again we have $\langle A \rangle = 0$, but now it is easy to check that

$$\sigma^2 \cong N^2 \frac{\langle \tanh[y^* + z^* + \Delta(0)] \rangle_0^2 + \langle \tanh[y^* - z^* + \Delta(0)] \rangle_0^2}{2} = \left(\frac{2Nz^*}{\Gamma} \right)^2$$

is proportional to N^2. Hence this is a much less efficient state. Fig. 3 shows the behavior of σ^2/N^2 as a function of Γ. The inset shows how Γ_c depends on the spread of initial conditions. We conclude that the more heterogeneous the initial condition is, the more efficient is the final state and the more the fixed point y^* is stable.

The transition from a state where $\sigma^2 \propto N$ to a state with $\sigma^2 \propto N^2$ is generic in the minority game, and it has been been discussed by several authors [Savit et al., 1999, Johnson et al., 1999a, Challet and Marsili, 1999, Cavagna et al., 1999].

6 Analytic approaches to the Minority Game

The first significant attempts to understand the Minority Game dynamics have focused on the derivation of continuum time dynamical equations. This has been

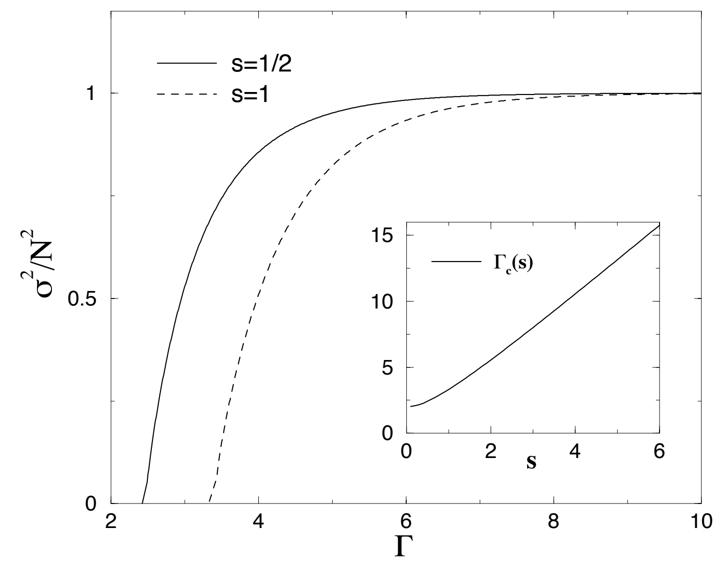

Figure 3: Global efficiency σ^2/N^2 as a function of Γ for two different sets of initial conditions: $\Delta_i(0)$ is drawn from a gaussian distribution with variance s^2. The full line corresponds to $s = 1/2$ whereas the dashed line is the result for $s = 1$. The inset reports the critical learning rate Γ_c as a function of the spread s of initial conditions.

the subject of debate for some time. First Cavagna *et al.* [Cavagna et al., 1999] proposed *stochastic* differential equations for time evolution of the probabilities

$$\pi_{i,s}(t) = \text{Prob}\{s_i(t) = s\}$$

which were found to be problematic in Ref. [Challet et al., 2000b] (but see also [Cavagna et al., 2000]) and later amended in Refs. [Garrahan et al., 2000, Sherrington et al., 2002]. On the other hand, Refs. [Challet et al., 2000c, Marsili et al., 2000] derived a *deterministic* dynamical equations by erroneously neglecting stochastic fluctuations. The asymptotic state of the dynamics was found to be related to the minima of H, which is a Lyapunov function of the deterministic dynamics. This opened the way to statistical mechanics of disordered systems because it relates the properties of the stationary state of the Minority Game to the ground state properties of a disordered spin model, which can be analysed in all details. It is remarkable that, in spite of neglecting fluctuations, this approach yields very precisely the behaviour of σ^2 and H with α, at least for $\alpha > \alpha_c$ for all values of Γ and for $\alpha < \alpha_c$ in the limit $\Gamma \to 0$. The reason for this coincidence was found in Ref. [Marsili and Challet, 2001a] which restores the stochastic term in the dynamics of Refs. [Challet et al., 2000c, Marsili et al., 2000] and provides a coherent picture of the Minority Game behaviour. In addition, the generating functional approach substantiated this approach [Coolen and Heimel, 2001].

It turns out that it is not necessary to derive continuum time equations in order to show that the minima of H describe several quantities in the stationary state of the Minority Game. Therefore we shall outline the theoretical developments introducing the continuum time limit only in a second stage.

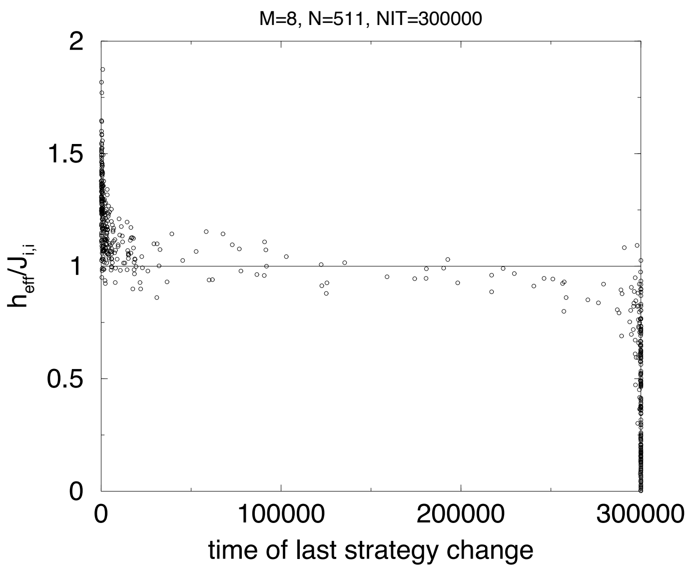

Figure 4: The condition for an agent to be frozen requires that the field $h_i = \overline{\Omega\xi_i} + \sum_{j\neq i}\overline{\xi_i\xi_j}m_j$ be larger in absolute value than the self-interaction term $J_{i,i} = \overline{\xi_i^2} \simeq 1/2$. Indeed $m_i = -\text{sign}\,[h_i + J_{i,i}m_i]$ has only a solution when $|h_i| > J_{i,i}$. Indeed plotting the time of the last change of strategy of agents versus the ratio $|h_i|/J_{i,i}$ we see that all those agents with $|h_i| > J_{i,i}$ soon freeze onto one strategic choice. The plot refers to a simulation with $M = 8$, $N = 511$, $S = 2$, 310^5 iterations. Here $h_{\text{eff}} = \overline{\Omega\xi_i}$.

We shall keep the discussion as simple as possible and refer the reader to the relevant original papers. Some of the results can also be derived with the generating functional method [Heimel and Coolen, 2001, Coolen and Heimel, 2001], as discussed in Section 6.6. Finally we shall comment on how the theory extends to several 'variations on the theme'.

6.1 Stationary state and minimal predictability

Taking the average of Eq. (20) in the stationary state one finds a dynamic equation for $\langle Y_i\rangle$. We look for solutions with $\langle Y_i\rangle \sim v_i t$. If we define

$$m_i \equiv \langle \tanh(Y_i)\rangle = \langle s_i\rangle, \tag{33}$$

for $t\to\infty$ we have

$$v_i = -\overline{\Omega\xi_i} - \sum_{j=1}^{N}\overline{\xi_i\xi_j}m_j. \tag{34}$$

Now if $v_i \neq 0$, then y_i diverges $\to \pm\infty$ and

$$m_i = \text{sign}\, v_i = -\text{sign}\left[\overline{\Omega\xi_i} + \sum_{j=1}^{N}\overline{\xi_i\xi_j}m_j\right] = \pm 1. \tag{35}$$

This means that agent i will always use the strategy $s_i(t) = m_i$ for t large enough, i.e. that he/she is *frozen*. Conversely, agent i is not frozen if $\langle Y_i \rangle$ is finite, which requires $v_i = 0$, i.e.:

$$0 = -v_i = \overline{\Omega\xi_i} + \sum_{j=1}^{N} \overline{\xi_i\xi_j} m_j. \tag{36}$$

The presence of the self-interaction term $\overline{\xi_i^2} m_i$ in Eq. (36) is crucial [Challet and Marsili, 1999] for the existence of non-frozen agents, as shown in Fig. 6.1. We call $\mathcal{F}$ the set of frozen agents; $\phi = |\mathcal{F}|/N$ is the fraction of frozen agents. Eqs. (35,36) are equivalent to the solution of the constrained minimization problem

$$\min_{\{m_i\}} H, \qquad m_i \in [-1, +1] \quad \forall i \tag{37}$$

where

$$H = \overline{\langle A \rangle^2} = \frac{1}{P} \sum_{\mu=1}^{P} \left[\Omega^\mu + \sum_{i=1}^{N} \xi_i^\mu m_i \right]^2 \tag{38}$$

This is easily shown by taking the first order derivatives of H with respect to m_i. Either $\partial H/\partial m_i = -2v_i = 0$ and then m_i takes a value in the interval $(-1, 1)$ (and agent i is not frozen) or $\partial H/\partial m_i = -2v_i \neq 0$ and then $m_i = \operatorname{sign} v_i$.

All quantities, such as H or ϕ, which can be expressed in terms of m_i can be computed if one can solve the "static" problem Eq. (37). We call m_i "average behaviour" of agent i for short. It is represented by a so-called soft-spin in the statistical mechanics formalism.[12]

The fact that the stationary state behaviour is related to the minima of the predictability H is a quite robust feature of Minority Games. This is also a natural result: Each agent is trying to predict the market outcome A with the limited capabilities — the strategies — at his disposal. The only possible stationary state is one where the outcome A is as unpredictable as possible.

The solution $\{m_i\}$ does not depend on Γ and, by Eq. (38), neither does H. On the contrary σ^2 cannot be expressed in terms of m_i only. Indeed

$$\sigma^2 = H + \sum_{i=1}^{N} \overline{\xi_i^2}(1 - m_i^2) + \sum_{i \neq j} \overline{\xi_i\xi_j} \langle (\tanh Y_i - m_i)(\tanh Y_j - m_j) \rangle \tag{39}$$

The last term depends on fluctuation of $\tanh Y_i$ around m_i. It only involves off-diagonal correlations across different agents $i \neq j$. An analysis of fluctuations in the stationary state is necessary in order to compute this last term and hence σ^2. This in turn will require the introduction of the continuum time limit. Let us discuss the properties of the solution to the problem (37). We anticipate that the effective theory which we shall develop later, and which is remarkably accurate, shows that the last term in Eq. (39) vanishes for $\alpha > \alpha_c$ for all values of Γ and for $\alpha \leq \alpha_c$ in the limit $\Gamma \to 0$. Accordingly, the solution of Eq. (37) will also allow us to compute $\sigma^2 \cong H + \sum_i \overline{\xi_i^2}(1 - m_i^2)$ in these cases.

[12] A soft spin, as opposed to a (hard) spin $s_i = \pm 1$, is a real number $m_i \in [-1, 1]$.

6.2 The statistical mechanics analysis of stationary states and phase transition

Let us come back to the minimization problem in Eq. (37). The statistical properties of the solutions to this problem can be accessed using techniques of statistical mechanics. Indeed, regarding H as the Hamiltonian of a system of soft spins $\{m_i\}$, the solution to Eq. (37) is given by the associated ground state properties. These are studied first introducing the partition function

$$Z(\beta, \Xi) = \mathrm{Tr}_m e^{-\beta H\{m_i, \Xi\}}$$

where β is an inverse temperature, $\Xi = \{a^\mu_{i,s}\}$ denotes the quenched disorder and Tr_m stands for the integral on m_i from -1 to $+1$, for all $i = 1, \dots, N$. This is nothing else than a generating function, from which all the statistical properties can be computed. In our case, since we are interested in the minimum of H, we shall take the limit $\beta \to \infty$ at the end of the calculus. Eq. (37) can be rewritten as:

$$\min_{\{m_i\}} H\{m_i, \Xi\} = -\lim_{\beta\to\infty} \frac{1}{\beta} \ln Z(\beta, \Xi). \tag{40}$$

Rather than in the solution of this problem for a particular game, that is for a given realization of the structure of interactions Ξ, we are interested in the generic properties which hold for 'typical' realizations of the game in the limit $N \to \infty$. These properties are called *self-averaging* because they hold for almost all realizations. In other words, in this limit, all the realizations of the game are characterized by the same statistical behaviour, i.e. the same values for all the relevant quantities.[13] In order to study these properties, we take the average over the disorder Ξ in Eq. (40). This eliminates the dependence on quenched disorder but leaves us with the problem of taking the average of the logarithm of a random variable Z, which is very difficult at best. However, the replica trick [Mézard et al., 1987]

$$\langle \ln Z\rangle_\Xi = \lim_{n\to 0} \frac{1}{n} \ln \langle Z^n\rangle_\Xi. \tag{41}$$

reduces the complexity of the calculus. In this equation, Z^n means that one replicates n times a given system, keeping the same disorder (strategies), but introducing a set of variables m_i for each replica c, denoted by $\{m_i^c\}$. One finds that (see e.g. the appendix of [Challet et al., 2000d]) the calculation of $\langle Z^n\rangle_\Xi$ can be reduced to an integral on a space of $n\times n$ matrices $\hat{Q}$ and $\hat{R}$, i.e.

$$\langle Z^n\rangle_\Xi \propto \int d\hat{Q} d\hat{R} e^{-\beta N n f(\hat{Q}, \hat{R})} \simeq e^{-\beta N n f(\hat{Q}^*, \hat{R}^*)} \tag{42}$$

where the saddle point integration was used in the last step (note that both β and N are large). Here Q is a matrix of *order parameters* and it has elements

$$Q_{a,b} = \frac{1}{N} \sum_{i=1}^{N} m_{i,c} m_{i,d} \tag{43}$$

[13] This also means that when the system's size increases, one has to average over less samples in numerical simulations for a given desired accuracy.

where the indices $_c$ and $_d$ label replicas. The matrix $\hat{R}$ is introduced in order to enforce Eq. (43) as a Lagrange multiplier.The free energy f is given by:

$$\begin{aligned} f(\hat{Q},\hat{r}) &= \frac{\alpha}{2n\beta}\text{Tr}\,\log\left[\hat{1}+\frac{\beta}{\alpha}\left(1+\hat{Q}\right)\right]+\frac{\alpha\beta}{2n}\sum_{c\le d} r_{c,d}Q_{c,d} \\ &- \frac{1}{n\beta}\log\left[\text{Tr}_{\,m} e^{\frac{\alpha\beta^2}{2}\sum_{c\le d} r_{c,d}m_c m_d}\right]. \end{aligned} \tag{44}$$

It is known that for Hamiltonian which are non-negative definite, such as $H = \langle A\rangle^2$, the matrices $\hat{Q}^*$ and $\hat{R}^*$ which dominate the integrals in the limit $\beta N\to\infty$ have the *replica symmetric* form

$$Q^*_{a,b} = q+(Q-q)\delta_{a,b}, \qquad R^*_{a,b} = r+(R-r)\delta_{a,b}. \tag{45}$$

Another consequence of the non-negativity of H is that it takes its minima on a connected set, which is either a point or a linear subspace[14]. In section 9 we shall see a case where the minimized function is no more positive definite. With the *ansatz* (45), we can compute the free energy f and then take the limit $n\to 0$.

$$\begin{aligned} f^{(RS)}(Q,q,R,r) &= \frac{\alpha}{2\beta}\log\left[1+\frac{\beta(Q-q)}{\alpha}\right] \\ &+ \frac{\alpha}{2}\frac{1+q}{\alpha+\beta(Q-q)}+\frac{\alpha\beta}{2}(RQ-rq) \\ &- \frac{1}{\beta}\left\langle \log\int_{-1}^{1} dm e^{-\beta V_z(m)}\right\rangle_z \end{aligned} \tag{46}$$

where

$$V_z(m) = -\frac{\alpha\beta(R-r)}{2}m^2-\sqrt{\alpha r}\,z\,m \tag{47}$$

and the average of the last term is defined as $\langle\ldots\rangle_z = \int_{-\infty}^{\infty} dz\ldots e^{-z^2/2}/\sqrt{2\pi}$ The last term of $f^{(RS)}$ looks like the free energy of a particle in the interval $[-1,1]$ with potential $V_z(m)$ where z plays the role of disorder. The parameters Q,q,R and r are finally found solving the saddle point equations $\frac{\partial f^{(RS)}}{\partial X}=0$ with $X = Q,q,R$ or r. The derivation of these equations is standard in disordered spin systems [Mézard et al., 1987] and are described in some more detail in Refs. [Marsili et al., 2000, Challet et al., 2000d].

The properties of the solution differ qualitatively according to whether $\alpha > \alpha_c$ or $\alpha < \alpha_c$ where

$$\alpha_c = 0.3374\ldots \qquad (S=2) \tag{48}$$

is the solution of the non-linear equation $\alpha_c = \text{erf}\left[\sqrt{|\log[\sqrt{\pi}(2-\alpha_c)|]}\right]$ where $\text{erf}\,(x)$ is the error function.

[14]A remarkable consequence of this will be discussed in section 6.4.

6.3 The asymmetric phase

For $\alpha > \alpha_c$ the solution to Eq. (37) is unique. In parametric form, we find,

$$\lim_{N\to\infty} \frac{1}{N}\langle \min_{\{m_i\}} H\{m_i\}\rangle_\Xi = \frac{1+Q}{2(1+\chi)^2} \tag{49}$$

where the parameters Q and χ are given by

$$Q(\zeta) = 1 - \sqrt{\frac{2}{\pi}}\frac{e^{-\zeta^2/2}}{\zeta} - \left(1-\frac{1}{\zeta^2}\right)\operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right) \tag{50}$$

$$\chi(\zeta) = \frac{\operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right)}{\alpha - \operatorname{erf}\left(\frac{\zeta}{\sqrt{2}}\right)} \tag{51}$$

where ζ is determined by

$$\alpha = [1+Q(\zeta)]\zeta^2, \tag{52}$$

for $\alpha > \alpha_c$[15].

The parameter $Q = \frac{1}{N}\sum_i m_i^2$ emerges in the calculation as an *order parameter*. It provides a measure of the degree of randomness of agents' behaviour: $Q = 1$ means that all agents stick to only one strategy ($m_i = \pm 1$ for all i) whereas $Q = 0$ means that they play at random ($m_i = 0$). A similar measure is given by the fraction ϕ of frozen agents, which is given, in parametric form, by

$$\phi = \lim_{N\to\infty}\frac{|\mathcal{F}|}{N} = \operatorname{erfc}(\zeta/\sqrt{2}). \tag{53}$$

A more detailed information on how agents play, is given by the full probability distribution of m_i:

$$\begin{aligned}\mathcal{P}(m) &= \lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\delta(m_i - m) \\ &= \frac{\phi}{2}\delta(m+1) + \frac{\phi}{2}\delta(m-1) + \frac{\zeta}{\sqrt{2\pi}}e^{-\zeta^2 m^2/2}.\end{aligned} \tag{54}$$

The quantity χ emerges instead as a *response function* in the statistical mechanics approach. More precisely, it is given by

$$\chi = \lim_{\beta\to\infty}\frac{\beta(Q-q)}{\alpha} = \lim_{\beta\to\infty}\frac{\beta}{\alpha N}\sum_{i=1}^{N}(m_{i,c} - m_{i,d})^2. \tag{55}$$

$Q-q$ measures the distance between two different replicas of the system, labelled by the indices $_c$ and $_d$ in the above equation. We can think of a replica as a realization of the stochastic process with given initial conditions. A finite value of χ means simply that two processes with different initial conditions converge, in the stationary state, to the same point in phase space $\{m_i\}$, i.e. that $q \to Q$ as $\beta \to \infty$ in the statistical mechanics formalism. This is what we expect to occur in an ergodic Markov process. We shall see that for $\alpha < \alpha_c$ the process

[15] In practice, Eqs. (50,51) and (52) describe the $\alpha > \alpha_c$ phase in a parametric form, for $\zeta > \zeta_c$ where ζ_c is the solution of $[1+Q(\zeta)]\zeta^2 = \operatorname{erf}(\zeta/sqrt2)$ (and $\alpha \to \alpha_c$ as $\zeta \to \zeta_c$).

is not ergodic, i.e. the stationary state depends on initial conditions. Indeed $\chi = \infty$ for $\alpha < \alpha_c$ as we shall see.

The behaviour of the solution as a function of α is the following: When $\alpha \gg 1$ agents behave nearly randomly $Q \approx 0$. A naive explanation is that the information encoded in μ is too complex and agents with their limited processing power are unable to detect significant patterns.

When α decreases agents manage to exploit information in a more efficient way: Hence Q (and ϕ) increases – implying a larger specialization in the population – and H decreases. Note that a decrease in α corresponds either to a decrease in the complexity P of the information, or to an increase in the number N of agents, and thereby of their collective information processing power. Hence at fixed P we find that as more and more agents join the game, the game's outcome becomes less and less predictable. This is a quite reasonable property for such complex adaptive systems.

Decreasing α we also find that χ increases more and more steeply. $\chi \sim (\alpha - \alpha_c)^{-1}$ diverges as $\alpha \to \alpha_c$ and correspondingly $H \sim (\alpha - \alpha_c)^2 \to 0$ (see Eq. 49). This singularity marks the location of a phase transition. For $\alpha > \alpha_c$ we have $H > 0$ which means that, given the information μ, the outcome $A(t)$ is probabilistically predictable, we call this an *asymmetric phase*. For $\alpha < \alpha_c$ we have $H = 0$ which means that for any μ the outcome $A(t)$ is symmetric. So we call this the symmetric phase.

6.4 The symmetric phase and dependence on initial conditions

It is worth to notice that the occurrence of a phase transition can be understood from a simple algebraic argument. Consider the set of $N(1-\phi)$ unfrozen agents, those who have $|m_i| < 1$. In the stationary state the corresponding variables m_i must satisfy the set of linear equations:

$$\overline{\langle A \rangle \xi_i} = \overline{\Omega \xi_i} + \sum_{j=1}^{N} \overline{\xi_i \xi_j} m_j = 0.$$

There are at most P independent equations in this set because that is the rank of the matrix $\overline{\xi_i \xi_j}$. Hence as long as $N(1-\phi) < P$ the solution is unique but when $N(1-\phi) \geq P$ there are more variables than equations and the solution is no more unique. Notice that in this case we have $\langle A | \mu \rangle = 0$ for all μ which means $H = 0$ for $N(1-\phi) \geq P$. We conclude that the critical threshold is given by:

$$\alpha_c = 1 - \phi(\alpha_c) \tag{56}$$

and that for $\alpha < \alpha_c$ the solution is no more unique, a fact which is at the origin of the divergence of χ in the whole $\alpha \leq \alpha_c$ phase. Note that the replica method confirms the validity of Eq (56).

The non-uniqueness of the stationary state implies that the properties of the Minority Game in the symmetric phase depend on the initial conditions. This was first observed in Ref. [Challet et al., 2000c] then confirmed numerically in Ref. [Garrahan et al., 2000]. Finally Ref. [Marsili and Challet, 2001a] showed that it is possible to characterise this dependence in the limit $\Gamma \to 0$

where the dynamics becomes deterministic. Similar conclusions extend qualitatively to the more complex stochastic dynamics ($\Gamma > 0$), as discussed in Ref. [Marsili and Challet, 2001a]. Before coming to that, let us mention that much insight can be gained on the dynamics in the symmetric phase by studying the limit $\alpha \to 0$ [Marsili, 2001, Marsili and Challet, 2001b].

In the $\alpha < \alpha_c$ phase the minimum of H is degenerate. In order to select a particular solution of $H = 0$ Ref. [Marsili and Challet, 2001a] adds a potential $\eta \sum_i (m_i - m_i^*)^2/2$ to the Hamiltonian H. This term lifts the degeneracy and selects the equilibrium close to m_i^* in the limit $\eta \to 0$. We refer to Ref. [Marsili and Challet, 2001a] for details and focus on the main results here.

Taking $m_i^* = m^* = 0$ describes symmetric initial conditions and increasing $m^* > 0$ gives asymmetric states that are reached when the initial scores are biased, that is $Y_i(0) \neq 0$. In this case the saddle point equations of the statistical mechanics approach reduce to:

$$Q = \int_{-\infty}^{\infty} Dz m_0^2(z) \tag{57}$$

$$\chi = \frac{1+\chi}{\sqrt{\alpha(1+Q)}} \int_{-\infty}^{\infty} Dz z m_0(z) \tag{58}$$

where $Dz = \frac{dz}{\sqrt{2\pi}} e^{-z^2/2}$ and $m_0(z) \in [-1, 1]$ is the value of m which minimises

$$V_z(m) = \frac{1}{2}m^2 - \sqrt{\frac{1+Q}{\alpha}} zm + \frac{1}{2}\eta(1+\chi)(m - m^*)^2. \tag{59}$$

There are two possible solutions: one with $\chi < \infty$ finite as $\eta \to 0$ which describes the $\alpha > \alpha_c$ phase. Note that, as long as χ remains finite, the last term of Eq. (59) vanishes when $\eta \to 0$, hence the dependence on 'initial conditions' m^* disappears.

The second solution has $\chi \sim 1/\eta$ which diverges as $\eta \to 0$, hence the last term of Eq. (59) has a finite limit. This solution describes the $\alpha < \alpha_c$ phase and can be expressed in parametric form in terms of two parameters z_0 and $\epsilon_0 = cm^*/(1 + c)$, where $c = \lim_{\eta \to 0} \eta\chi$. Eq. (57) with

$$m_0(z) = \begin{cases} -1 & \text{if } z \le -z_0 - \epsilon_0 \\ \frac{z+\epsilon_0}{z_0} & \text{if } -z_0 - \epsilon_0 < z < z_0 - \epsilon_0 \\ 1 & \text{if } z \ge z_0 - \epsilon_0 \end{cases}$$

gives $Q(z_0, \epsilon_0)$ and Eq. (58), which for $\chi \to \infty$ reads $\sqrt{\alpha(1+Q)} = \int Dz z m_0(z)$, then gives $\alpha(z_0, \epsilon_0)$. With $\epsilon_0 \neq 0$, i.e. $m^* \neq 0$, one finds solutions with a non-zero 'magnetization' $M = \sum_i m_i/N$. This quantity is particularly meaningful, in this context, because it measures the overlap of the behaviour of agents in the stationary state with their a priori preferred strategies

$$M \equiv \int_{-\infty}^{\infty} Dz\, m_0(z) = \lim_{t \to \infty} \frac{1}{N} \sum_{i=1}^{N} \langle s_i(t) s_i(0) \rangle. \tag{60}$$

Which stationary state is reached from a particular initial condition is a quite complex issue which requires the integration of the dynamics. However, the relation between Q and M derived analytically from Eqs. (57) and (60), which is

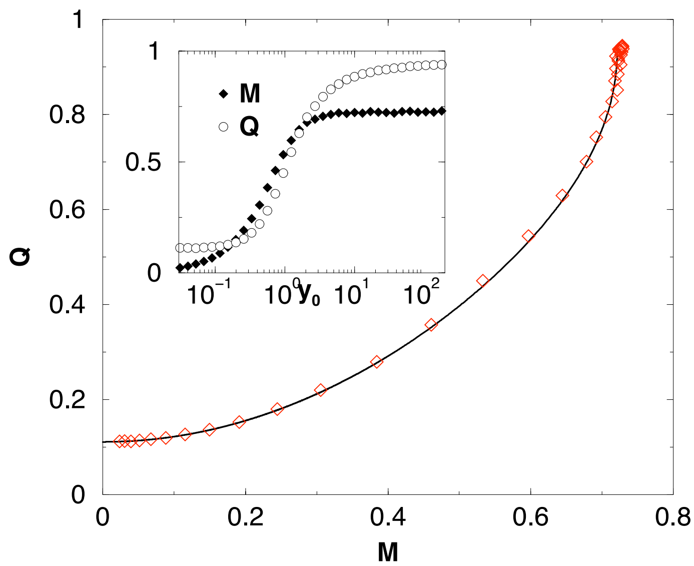

Figure 5: Relation between Q and M, for $\alpha = 0.1$, derived from analytic calculation (full line) and from numerical simulations of the MG with different initial conditions y_0 ($\diamond$, $P = 32$, $N = 320$, $\Gamma = 0.1$). The inset shows the dependence of Q and M from the initial condition y_0.

a sort of 'equation of state', can easily be checked by numerical simulations of the Minority Game. Figure 5, from Ref. [Marsili and Challet, 2001a], shows that the self-overlap Q and the magnetization M computed in numerical simulations with initial conditions $y_i(0) = y_0$ for all i, perfectly match the analytic results. The inset of this figure shows how the final magnetization M and the self-overlap Q depend on the asymmetry y_0 of initial conditions.

In order to show the variability of results with initial conditions, Fig. 6 plots σ^2/N both for symmetric ($y_0 = 0$) and for maximally asymmetric initial conditions ($y_0 \to \infty$) in the limit $\Gamma \to 0$. The inset shows the behaviour of Q and M in the maximally asymmetric state.

Remarkably we find that σ^2/N vanishes as $\alpha \to 0$ in the maximally asymmetric state[16], in agreement with the results of Ref. [Heimel and Coolen, 2001]. This means that, at fixed P, as N increases the fluctuation σ^2 remains constant. This contrast with what happens in the $y_0 = 0$ state, for $\Gamma \ll 1$, where σ^2 increases linearly with N, and with the case $\Gamma \to \infty$ where $\sigma^2 \propto N^2$ [Savit et al., 1999].

6.5 The continuum time limit and fluctuations in the stationary state

The study of the off-diagonal correlations $\langle s_i s_j \rangle$ requires a more refined approach. Let us go back to Eq. (20). The study of the dynamics of the Minority Game starts from three key observations:

[16] The relation is almost indistinguishable from a linear law $\sigma^2 = cN\alpha$ but higher order terms exist.

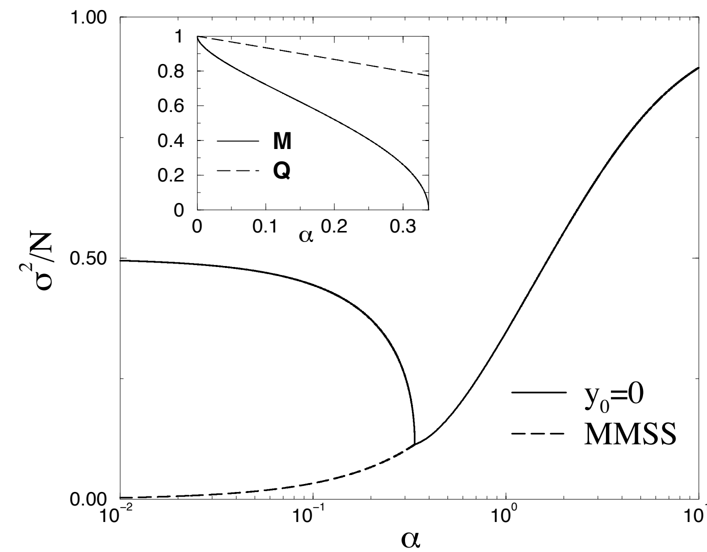

Figure 6: σ^2/N for the Minority Game with initial conditions $y_0 = 0$ (full line) and $y_0 \to \infty$ (dashed line). The inset reports the behaviour of M and Q for $y_0 \to \infty$.

1. the scaling $\sigma^2 \sim N$, at fixed α, suggests that typically $A(t) \sim \sqrt{N}$. Hence time increments of $U_{i,s}(t)$, in Eq. (16) are small (i.e. of order $1/\sqrt{N}$);

2. as shown in Ref. [Challet et al., 2000b], characteristic times of the dynamics are proportional to N. Naively this is because agents need to 'test' their strategies against all $P = \alpha N$ values of μ, which requires of order P time steps.

3. characteristic times in the learning dynamics are also inversely proportional to Γ. The process takes a time of order $1/\Gamma$ to 'learn' a perturbation $\delta U_{i,s}$, i.e. to translate it into a change in the choice probabilities Eq. (13) of the same order. From this perspective Γ appears as a 'learning rate' rather than the inverse of a temperature.

The second and last observation implies that one needs to study the dynamics in the rescaled time

$$\tau = \frac{\Gamma t}{N}$$

and to introduce continuum time variables $y_i(\tau) = Y_i(t)$. The key point is that if $N \gg \Gamma$, a small time increment $d\tau \ll 1$ corresponds to a large number $\Delta t = N d\tau/\Gamma \gg 1$ of time-steps. The corresponding change $dy_i(\tau) = Y_i(t+\Delta t) - Y_i(t)$, being the sum of Δt stochastic increments, can be estimated quite precisely by the Central Limit Theorem when $\Delta t \to \infty$. Having taken the thermodynamic limit $N \to \infty$, one can take the continuum time limit $d\tau \to 0$. Note that Γ needs to be finite in this process, and the limit $\Gamma \to \infty$ can be taken at the end. In practice, numerical simulations show that the limit in which the limits $N \to \infty$ and $\Gamma \to \infty$ are taken does not matter.

We refer the interested reader to Ref. [Marsili and Challet, 2001a] for a detailed account of the derivation and jump directly to the resulting dynamical equations:

$$\frac{dy_i}{d\tau} = -\overline{\Omega\xi_i} - \sum_{j=1}^{N} \overline{\xi_i\xi_j} \tanh(y_j) + \zeta_i \tag{61}$$

where $\zeta_i(\tau)$ is white noise with $\langle\zeta_i(\tau)\rangle = 0$ and

$$\langle\zeta_i(\tau)\zeta_j(\tau')\rangle = \frac{\Gamma}{N}\overline{\langle A^2\rangle_y \xi_i\xi_j}\delta(\tau - \tau') \tag{62}$$

The average $\langle\ldots\rangle_y$ in Eq. (62) is taken over the instantaneous probabilities $\text{Prob}\{s_i(t) = s\} = (1+s\tanh y_i)/2$ in Eq. (19) of $s_i(t)$. In other words, the noise covariance depends in a non-linear and complex way on the dynamical variables $y_i(\tau)$. Hence Eqs. (61,61) are complex non-linear stochastic differential equation with a time dependent noise term. They are exact in the limit $N \to \infty$ with Γ finite. This conclusion has been confirmed by the more elaborate generating functional approach of Ref. [Heimel and Coolen, 2001, Coolen and Heimel, 2001] (see section 6.6).

A peculiar feature of these equations is that the noise strength in Eq. (62) is itself proportional to $\langle A^2\rangle_y \approx \sigma^2$. This feedback effect is quite natural in hindsight: Each agent faces an uncertainty which is large when the volatility $\langle A^2\rangle_y$ is large.

Quite remarkably, the stochastic force is proportional to Γ and that's the only place where Γ appears explicitly. Hence Γ tunes the strength of stochastic fluctuations in much the same way as temperature does for thermal fluctuations in statistical mechanics. It is significant that Γ, which is introduced as the inverse of an *individual* temperature in the definition of the model, actually turns out to play *collectively* a role quite similar to that of global temperature. This similarity will appear even more evident below.

The analysis of the stochastic dynamics is made complex by the dependence on $y_i(\tau)$, and hence on time, of Eq. (62). However this time dependence comes through the volatility $\langle A^2|\mu\rangle_y/N$ which is self-averaging unless collective fluctuations of the variables $y_i(\tau)$ arise. Ref. [Challet and Marsili, 2003a] expands further on this argument, showing that collective fluctuations which can sustain time dependent volatility fluctuations only arise close to the critical point and for finite size systems. Away from it, the feedback arising from time dependent volatility can be neglected assuming that

$$\overline{\langle A^2\rangle_y \xi_i\xi_j} \approx \overline{\langle A^2\rangle_y}\,\overline{\xi_i\xi_j} \approx \sigma^2\overline{\xi_i\xi_j}.$$

This greatly simplifies our task by replacing Eq. (62) with

$$\langle\zeta_i(t)\zeta_j(t')\rangle \cong 2T\overline{\xi_i\xi_j}\delta(t - t'), \quad \text{with } T = \frac{\Gamma\sigma^2}{2N} \tag{63}$$

One important point is that in going from Eq. (62) to (63) we pass from an exact to an effective theory. Note indeed that the temperature T in Eq. (63) depends on σ^2 which in its turn depends on the fluctuations of $y_i(\tau)$. So the theory becomes a self-consistent one. The comparison of its predictions with numerical simulations will provide a check of its validity.

Let us imagine we have solved the problem of computing m_i in the stationary state[17] and let us address the problem of computing the fluctuations of y_i. We briefly review the main steps of the analysis in Ref. [Marsili and Challet, 2001a] and refer the interested reader to the original paper for more details.

Using the stationary condition Eq. (36) and Eq. (63), we can write the Fokker-Planck equation for the probability distribution $P_u(y_i, i \notin \mathcal{F})$ of unfrozen agents. This satisfies a sort of fluctuation dissipation theorem: indeed both the deterministic term of Eq. (61) and the noise covariance Eq. (63) are proportional to the matrix $J_{i,j} = \overline{\xi_i \xi_j}$. This makes it possible to find a solution in the stationary state, which reads

$$P_u \propto \mathcal{P}_{y(0)} \exp\left\{-\frac{1}{T}\sum_{j\notin\mathcal{F}}[\log\cosh y_j - m_j y_j]\right\} \tag{64}$$

where

$$\mathcal{P}_{y(0)} \equiv \prod_{\mu=1}^{P}\int_{-\infty}^{\infty} dc^\mu \prod_{i=1}^{N}\delta\left[y_i - y_i(0) - \sum_{\mu=1}^{P} c^\mu \xi_i^\mu\right] \tag{65}$$

is a projector which imposes the constraint that the states $|y(t)\rangle = \{y_i(t)\}_{i=1}^N$ which are dynamically accessible must lie on the linear space spanned by the vectors $|\xi^\mu\rangle$ which contains the initial condition $|y(0)\rangle$. Note that Eq. (64) has the form of a Boltzmann distribution with temperature T, which is proportional to Γ (see Eq. 63).

Using the distribution Eq. (64), we can compute σ^2 from Eq. (39). In principle the third term of Eq. (39)

$$\Sigma(T) = \sum_{i\neq j}\overline{\xi_i\xi_j}\langle(\tanh y_i - m_i)(\tanh y_j - m_j)\rangle_T$$

(where $\langle\ldots\rangle_T$ stands for averages over P_u) depends on T which in its turn depends on σ^2 (see Eq. 63). Hence the stationary state is the solution of a self-consistent problem.

For $\alpha > \alpha_c$ the number $N - |\mathcal{F}| \equiv N(1-\phi)$ of unfrozen agents is less than P and the constraint is ineffective, i.e. $\mathcal{P}_{y(0)} \equiv 1$. Hence the dependence on initial conditions $y_i(0)$ drops out and the probability distribution of y_i factorises over i. As a consequence $\Sigma(T) \equiv 0$ vanishes identically. We conclude that, for $\alpha > \alpha_c$, $\sigma^2 = H + \sum_i \overline{\xi_i^2}(1-m_i^2)$ only depends on m_i and is independent of Γ. This is confirmed by numerical simulations to a remarkable accuracy: Fig. 7 shows that σ^2 stays constant when Γ varies over four decades.

When $\alpha < \alpha_c$, on the other hand, the constraint cannot be integrated out and the stationary distribution depends on the initial conditions. In addition $\mathcal{P}_{y(0)}$ also introduces a correlation in the fluctuations of y_i across the agents. This leads to a non vanishing contribution $\Sigma(T)$ in Eq. (39). Hence σ^2 turns out to be the solution of the self-consistent equation:

$$\sigma^2(\Gamma) = H + \sum_{i=1}^{N}\overline{\xi_i^2}(1-m_i^2) + \Sigma\left(\frac{\Gamma\sigma^2(\Gamma)}{2N}\right). \tag{66}$$

[17] We remark that for $\alpha > \alpha_c$ the solution of $\min H$ is unique and hence m_i depend only on the realization of disorder Ξ. For $\alpha < \alpha_c$ the solution is not unique, hence we also have the problem of finding which solution the dynamics selects, depending on the initial conditions.

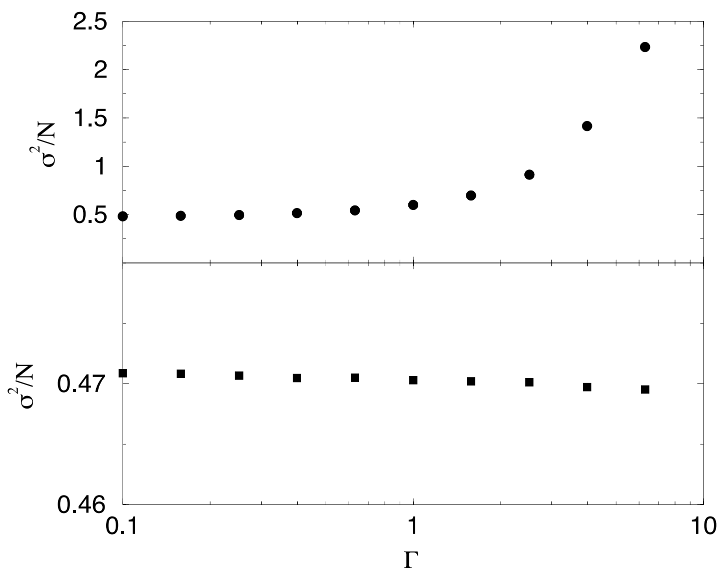

Figure 7: σ^2/N as a function of Γ for a given realisation of the game with initial conditions $y(0) = 0$; $\alpha \simeq 0.1 < \alpha_c$ (upper panel, $P = 10$, $N = 101$) and $\alpha = 1.5 > \alpha_c$ (lower panel, $P = 30$, $N = 45$). Average over $10000P$ iterations.

Ref. [Marsili and Challet, 2001a] solved these equations using Monte Carlo methods to sample the distribution P_u of Eq. (64). The results, reported in Fig. 8, agree perfectly with numerical simulations of the MG.

It is possible to solve Eq. (66) to leading order in $\Gamma \ll 1$. Ref. [Marsili and Challet, 2001a] shows that

$$\frac{\sigma^2}{N} \cong \frac{1-Q}{2}\left[1 + \frac{1-Q+\alpha(1-3Q)}{4}\Gamma + O(\Gamma^2)\right]. \tag{67}$$

which agrees very well with numerical simulations (see the inset of Fig. 8). Ref. [Marsili and Challet, 2001a] also shows that a simple argument allows to understand the origin of the behaviour $\sigma^2/N \sim 1/\alpha$ for $\Gamma \gg 1$, first discussed in Ref. [Savit et al., 1999]. It must be observed that a correlated behaviour of agents was already hinted at by Johnson and coworkers [Johnson et al., 1999a, Hart et al., 2001], who put this effect in relation with crowd effects in financial markets.

In summary the dependence on initial conditions, cross-correlations in the behaviour of agents and dependence of aggregate fluctuations on the learning rate are intimately related in a chain of consequences. This is a remarkable and entirely novel scenario in statistical physics. These results are derived under the approximation of Eq. (63) but are fully confirmed by numerical simulations. This suggests that this approximation may be exact in the limit $N \to \infty$ at least far from the critical point α_c.

A quantitative study of correlated fluctuations close to the critical point has been carried out in [Challet and Marsili, 2003a] on the basis of a simple argument: The time dependence of the volatility becomes relevant when

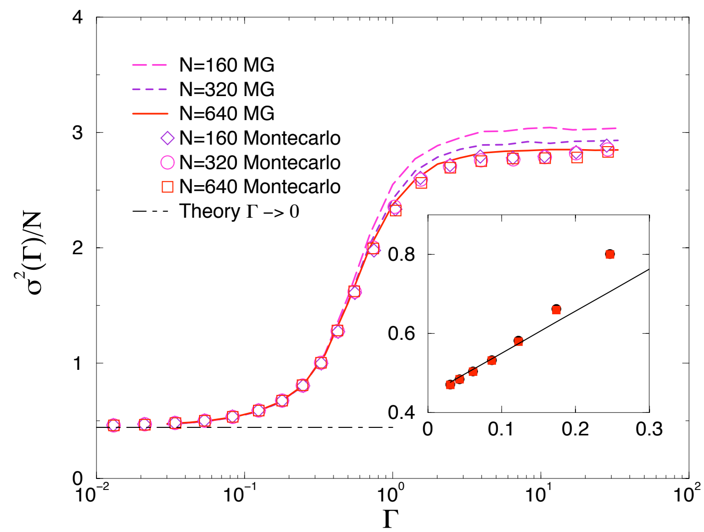

Figure 8: Global efficiency σ^2/N versus Γ for $\alpha = 0.1 < \alpha_c$ and different system sizes. Lines refer to direct simulations of the Minority Game with $N = 160, 320$ and 640. Finite size effect for $\Gamma \gg 1$ are evident. Symbols refer instead to the solution of the self-consistent equation (66) for the same system sizes. For both methods and all values of N, σ^2 is averaged over 100 realizations of the disorder. In the inset, the theoretical prediction Eq. (67) on the leading behaviour of σ^2/N for $\Gamma \ll 1$ (solid line) is tested against numerical simulations of the Minority Game (points) for the same values of N.

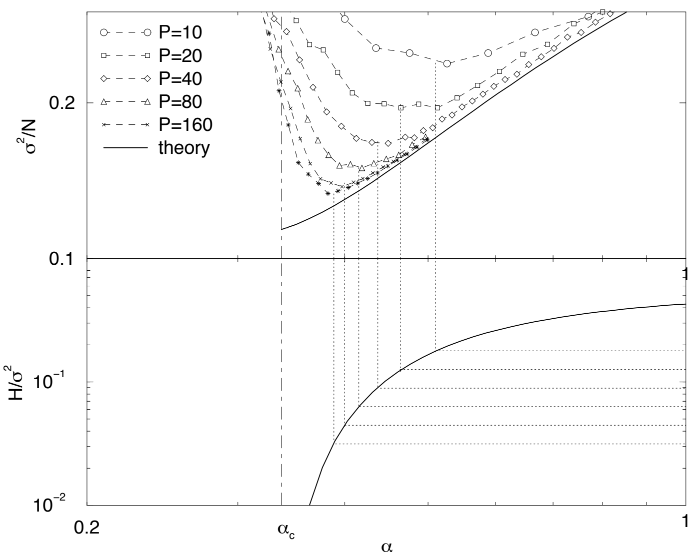

Figure 9: Graphic construction to show that deviations from theoretical result occur inside the critical region define by Eq. (68) ($K = 0.39$ in the figure).

the deterministic part of the dynamics (see Eq. (61)) is small. Then correlated fluctuations between y_i and y_j can be sustained by fluctuation in $\langle A^2|\mu\rangle$. The criteria for this to occur is, in the present model (see Eq. (6) in paper [Challet and Marsili, 2003a]), $H \ll \sigma^2/\sqrt{P}$. Indeed Fig. 9 shows that the point where

$$\frac{H}{\sigma^2} \simeq \frac{K}{\sqrt{P}} \tag{68}$$

with $K \approx 0.39$ a constant, determine quite precisely the location where numerical simulations deviate considerably from the theoretical results. Since $H \propto (\alpha - \alpha_c)^2$, we argue that the critical region where we expect anomalous fluctuations has a size which vanishes as $|\alpha - \alpha_c| \sim N^{-1/4}$ when $N, P \to \infty$.

We expect that Eq. (63) is a good approximation outside the critical region, whereas for $|\alpha - \alpha_c| \ll N^{1/4}$ one needs to take explicitly into account the time dependence of volatility.

6.6 The Generating functional approach

A different approach, based on the generating functional has been proposed to study the Minority Game [Heimel and Coolen, 2001, Coolen and Heimel, 2001, Heimel et al., 2001]. This method is dynamical in essence. Taking the average of dynamical equations over the disorder, it yields an effective dynamical theory. We include Ref. [Heimel and Coolen, 2001], which deals with the Minority Game with $\Gamma \to \infty$, in the list of reprints. This gives a quite detailed account of the method. We refer the interested reader to it while giving a brief account of the method and results in what follows.

This approach was proposed in Ref. [Heimel and Coolen, 2001] for a *batch* version of the Minority Game in which $\Gamma = \infty$ and the agents revise their choices $s_i(t)$ only every P time steps[18]. Later the analysis was extended to the standard – called on-line – Minority Game and later to $\Gamma < \infty$ and to a different choice rule [Coolen and Heimel, 2001, Heimel et al., 2001].[19]

The idea is to write down a dynamic generating functional as a path integrals over the exact time evolution of a given configuration of scores $\{U_{i,s}\}$. After taking the average over disorder, one finds that the generating functional is dominated (in the saddle point sense) by a single "representative" stochastic process $q(t)$ – Eqs. (62) to (66) of Ref. [Heimel and Coolen, 2001]. These are quite complex self-consistent equations: Indeed the drift and diffusion term of the process $q(t)$ depend in a non-linear way on the correlation and response functions

$$C(t,t') = \langle s(t)s(t')\rangle, \qquad G(t,t') = \frac{\delta\langle s(t)\rangle}{\delta h(t')}$$

of the process $s(t)$. Here $h(t)$ is an auxiliary external field which is added to the dynamics in order to probe its response. Furthermore the integration on the quenched disorder induces a long term memory in the process.

A virtue of this approach has been to clarify several issues on the continuum time limit[Cavagna et al., 1999, Challet et al., 2000c, Challet et al., 2000b, Cavagna et al., 2000]. In particular it has shown that characteristic times of the dynamics do indeed scale with N.[20] Furthermore it has shown that an approach based on Fokker-Planck equation (which neglects higher order terms in the Kramers-Moyal expansion) – such as that of previous sections [Marsili and Challet, 2001a] – is indeed correct. It also gives an exact derivation of the drift and diffusion terms in the continuum time description.

A drawback of the approach is that the resulting equations for $C(t,t')$ and $G(t,t')$ are too complex to be solved exactly and results are available only in limit cases or for some quantities. Assuming time translation invariance – $C(t,t+\tau) = C(\tau)$ and $G(t,t+\tau) = G(\tau)$ for $t \to \infty$ – a finite integrated response

$$\lim_{t\to\infty} \int dt' G(t,t') = \chi < \infty$$

and weak long term memory [Eq. (69) in Ref. [Heimel and Coolen, 2001]], Coolen and Heimel were able to re-derive the equations of the replica approach[21] and the phase transition at $\alpha_c = 0.3374\ldots$.

It is quite interesting that the generating function approach gives a dynamical interpretation of the quantity χ as the integrated response to an infinitesimal perturbation. Then $\chi = \infty$ implies that the system 'remembers' forever a perturbation. In particular this means that the stationary state depends on initial

[18]Put differently agents process in a batch all information patterns μ and react to their cumulative effect.

[19]Note that as long as the average behaviour m_i is concerned, this model is identical to the on-line case. Indeed Eqs. (34) are still valid. However, σ^2 differs from the online version of the game.

[20]In the batch version of the game, the time is *de facto* rescaled by a factor P.

[21]Eqs. (76,77) and (78) of Ref. [Heimel and Coolen, 2001], with $c \equiv C(\tau \to \infty) = q = Q$ and $y = \zeta$ become equivalent to the saddle point equations (50,51) and (52) discussed above. In addition Eq. (84) of Ref. [Heimel and Coolen, 2001] is exactly the expression of H in Eq. (49).

conditions. However when $\chi = \infty$, for $\alpha < \alpha_c$, the self-consistent dynamic equations are much more difficult to analyze and have not yet been solved .

The generating functional approach is a quite promising tool for these systems which has also been successfully extended to more complex cases [Heimel and Martino, 2001]. However it does not yet give a clear picture of the interplay between initial conditions, correlations and fluctuations – including the independence of σ^2 on Γ for $\alpha > \alpha_c$ – such as the one given by the approach outlined in the previous sections. On the other hand, the exact representative agent process that this method allows to derive makes it possible to carry out accurate numerical simulations of the system in the thermodynamic limit [Eissfeller and Opper, 1992].

7 Extensions

The analytic approaches discussed so far extend, in a more or less straightforward way to more complex models. Many of them lead to a quite similar generic picture characterised by a similar phase diagram. In many cases, the stationary state is related to the minimum of a functional which can be studied exactly within the replica symmetric approximation. We list here some of these extensions:

- The Minority Game with endogenous information has been studied in Ref. [Challet and Marsili, 2000]. It turns out that for $\alpha > \alpha_c$ the behaviour of the Minority Game with endogenous information (Eq. 6) slightly differs from that under exogenous information (Eq. 7). The correction can be quantified within the analytic approach of Refs. [Challet et al., 2000c, Marsili and Challet, 2001a], as shown in Ref. [Challet and Marsili, 2000]. In brief, the dynamics of $A(t)$ induces a dynamics on $\mu(t)$ according to Eq. (6). $\mu(t)$ is a diffusion process on a particular graph, called after De Bruijn. The diffusion process acquires a drift for $\alpha > \alpha_c$ because $\langle A|\mu\rangle \neq 0$. This results in the fact that some value of μ arise more frequently than others, i.e. the stationary state distribution ρ^μ of $\mu(t)$ is not uniform, as in the case of exogenous information ($\rho^\mu = 1/P$). These considerations can be cast into a self-consistent theory which approximately accounts for the effects of endogenous histories for $\alpha > \alpha_c$.

  For $\alpha < \alpha_c$ instead the stationary state is uniform ($\rho^\mu = 1/P$) because there is no bias ($\langle A|\mu\rangle = 0$) but correlation functions exhibit an oscillatory behaviour (Fig. 2 in Ref. [Challet and Marsili, 1999]) under endogenous information dynamics which does not arise with exogenous information. Refs. [P. Jefferies, 2001, Hart et al., 2002, Hui et al., 1999, Challet et al., 2000d, Metzler, 2002] also discuss the issue of endogenous versus exogenous information in different variants of the Minority Game. It should be noted that there are cases where endogeneous information makes sense, for instance in models of prediction, bubbles and crashes [Lamper et al., 2002, Giardina and Bouchaud, 2002].

  Endogenous information leads to a radically different results with respect to the endogenous case when agents behave in a deterministic way (i.e. when they are all frozen). Then the induced dynamics of $\mu(t)$ is also deterministic and it locks into periodic orbits of period $\sim \sqrt{P}$ on the De Bruijn graph. The majority of the information patterns are

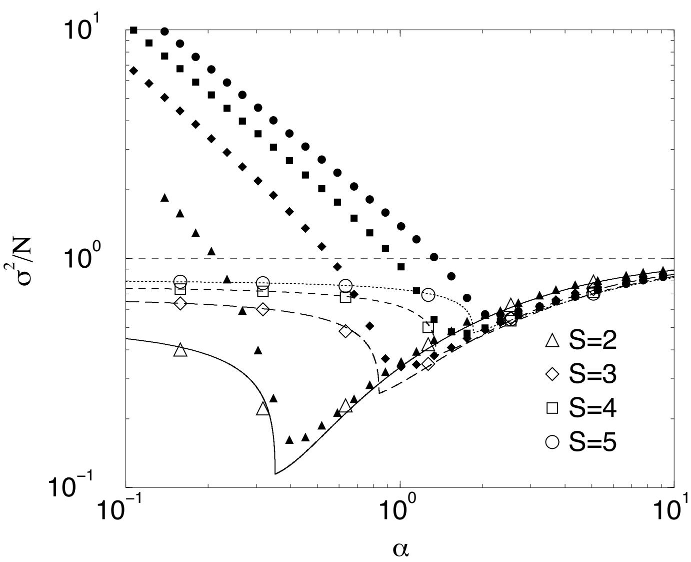

Figure 10: Global efficiency σ^2/N as a function of α for $S = 2, 3, 4$ and 5 from numerical simulations of the minority game with $N = 101$ agents and $\Gamma_i = \infty$ (averages were taken after $100P$ time steps), averaged over 100 realizations of $\vec{a}_i$ (small full symbols), from numerical minimization of H (large open symbols) and from the theoretical calculation (lines) with $U_{i,s}(0) = 0$.

never generated by the system. Such a situation has been discussed in Refs. [Johnson et al., 1999b, Challet et al., 2000c, Marsili et al., 2000, Challet and Marsili, 2003b].

- This whole approach can be generalised to $S > 2$ strategies [Marsili et al., 2000]. Fig. 10 shows that σ^2/N increases with S toward the random limit $\sigma^2/N = 1$ and the phase transition $\alpha_c(S) = \alpha_c(S = 2) + S/2 - 1$ moves linearly to higher values. Giving agents more resources leads generally to a smaller efficiency, because their strategy sets are less internally correlated.

- The approach has been also generalised to include a fraction of deterministic agents – i.e. agents with $S = 1$ strategy – and totally random agents. These extensions are discussed in Ref. [Challet et al., 2000d] and in the next chapter.

- Cavagna *et al.* [Cavagna et al., 1999] have proposed a generalization of the Minority Game where each agent contributes an action

$$a_i(t) = \sum_{\mu=1}^{P} a^{\mu}_{i,s_i} \eta^{\mu}(t) \tag{69}$$

where $\eta^\mu(t)$ is white noise (i.e. a Gaussian variable with zero average and unit variance, independent for each μ and t). In Ref. [Cavagna et al., 1999] trading strategies a^{μ}_{i,s_i} are "continuous", i.e. they are drawn from a continuous distribution – a Gaussian – rather than from the bimodal. The idea is that the exogenous process which drives the market, or the news arrival process, is a P dimensional vector $\vec{\eta}(t)$. Agents respond to it with linear strategies, which are also P dimensional vectors $\vec{a}_{i,s}$. This model reduces to the previous one if we assume that $\vec{\eta}(t) = (\ldots, 0, 1, 0, \ldots)$ can only lay along one of the P components of the orthogonal basis and that $a_{i,s}^{\mu}$ is drawn from the bimodal distribution. In general the vector $\vec{\eta}$ probes the performance of the strategies with respect to all informations (i.e. components) μ at the same time. Challet *et al.* [Challet et al., 2000b] have shown that this model has the same collective behaviour as the one discussed above. Indeed the equations which describe the stationary state of this model has the same deterministic term as Eq. (61).

- The agents can be allowed not to play by endowing them with a so-called $0-$strategy which prescribes not to trade, whatever μ. These models are characterised by a behaviour similar to that discussed here [Challet et al., 2001, Challet and Marsili, 2003a]. Even if, in general, agents do not play at all times, they all update their (virtual) payoffs simultaneously observing the market. In other words, the model is still a fully connected, mean field model. This case is dealt with in detail in the following section.

- The case of agents entering the market with heterogeneous weights $w_i = |a_{i,s}^{\mu}| \neq 1$ – modelling a population where some are richer or more influential than others – has been dealt with in Ref. [Challet et al., 2000a].

- Agents with a given degree of correlation between their two strategies [Challet et al., 2000d, Sherrington et al., 2002, Sherrington and Galla, 2003] also exhibit the same generic behaviour, with α_c depending on the correlation coefficient. Ref [Challet et al., 2000d] discusses this case within the replica approach, while ref [Sherrington and Galla, 2003] employ generating functionals.

- Agents playing on different frequencies have been discussed in Ref. [Marsili and Piai, 2002]. This amounts in assuming that agent i plays only for a fraction f_i of information patterns μ and otherwise shedoes not play ($a_{i,s}^{\mu} = 0$). Even for broad distribution of frequencies f_i across the population of agents the phase transition persists, with α_c which depends on the distribution of f_i. Interestingly it turns out that frequent players have a smaller probability of being frozen than those who play rarely.

- Information patterns with widely spread frequencies have been discussed in Ref. [Challet and Marsili, 2000, Marsili and Piai, 2002]. Again the phase transition persists and α_c depends on the distribution of the frequencies ρ^{μ} with which information patterns occur. It turns out that in the asymmetric phase, more frequent patterns are typically less predictable. More precisely if μ occurs with a probability $\rho^{\mu} = \tau^{\mu}/P$

  $$|\langle A|\mu\rangle| \propto \frac{1}{1+\chi\tau^{\mu}}$$

  which also suggests that χ is the inverse of the characteristic frequency above which the market is unpredictable ($\langle A|\mu\rangle \approx 0$ for $\tau^{\mu} \gg \chi^{-1}$).

- Minority games with non-linear payoffs $-a_i G(A)$ have been discussed in some details in Ref. [Li et al., 2000] on the basis of numerical simulations and in Ref. [Challet et al., 2000d, Marsili and Challet, 2001a] analytically. Ref [Marsili and Challet, 2001a] shows that, within a self-consistent time-independent volatility approximation, it is again possible to derive a function $\mathcal{H}$ which is minimized by agents in the stationary state. For $G(A) = -G(-A)$ and $G(A) \simeq gA + O(A^3)$ for $A \ll 1$, it is possible to argue that the location of the phase transition does not depend on $G(A)$ [Challet, 2003].

- The statistical mechanics approach makes it clear that what matters only the first two moments of the distribution of $a_{i,s}^{\mu}$. Our results stay exactly the same for all distributions of $a_{i,s}^{\mu}$ with zero average and unit variance. The case of a non zero average has been dealt with in Ref. [Challet and Marsili, 2003b].

- Ref. [Martino et al., 2003] has shown that an additive noise in the payoff dynamics does not change the phase diagram, but it affects the fluctuations σ^2. In particular Ref. [Martino et al., 2003] considers the case

$$U_{i,s}(t+1) = U_{i,s}(t) - a_{i,s}^{\mu} A(t) + (1 - \delta_{s,s_i(t)})\zeta_{i,s}(t) \qquad (70)$$

  where the payoffs of strategies $s \neq s_i(t)$ which have not been played are affected by a 'measurement' noise $\zeta_{i,s}(t)$ with zero average and variance Δ. While this term does not modify the conclusion that the stationary state behaviour is related to the minimum of H, it changes the fluctuation properties. The new term $\zeta_{i,s}(t)$ removes the degeneracy of stationary states for $\alpha < \alpha_c$ thus affecting the fluctuations. Remarkably the more noisy the estimate of payoffs of agents, the less noisy the aggregate behaviour, i.e. σ^2 turns out to be a *decreasing* function of Δ.

Further modifications which lead to a qualitatively different behaviour (characterised by replica symmetry breaking) or to a qualitatively different model with a similar behaviour will be discussed in the following two chapters.

8 Recovering financial markets' stylized facts in Minority Games

In this section, we first introduce the simplest possible Grand Canonical Minority Game (GCMG) which reproduces the main stylized facts, i.e. fat tails and volatility clustering. Then we present the analytic solution of this model in the relevant thermodynamic limit. It shows that the behavior of GCMG, in this limit, exhibits Gaussian fluctuation for all parameter values but on a line of critical points which marks a discontinuous phase transition. For finite size systems, numerical simulations reveal that stylized facts emerge close to the transition line, but they abruptly disappear as the system size increases. Remarkably, the vanishing of stylized facts when the system's size increases also occurs in a variety of models of financial markets. We present a theory of finite size effects which is fully confirmed by numerical simulations. This allows us to conclude that anomalous fluctuations are properties of the critical point in GCMG. The phase transition is quite unique as it mixes features which are typical of first order phase transitions – as discontinuities and phase coexistence – and of second order phase transitions – such as the divergence of correlation volumes and finite size effects.

In the market described by the Minority Game [Challet and Zhang, 1997], agents $i = 1, \cdots, N$ submit a bid $b_i(t)$ to the market in every period $t = 1, 2, \ldots$. Agents whose bid has the opposite sign of the total bid $A(t) = \sum_i b_i(t)$, win whereas the others lose. The bids of agents depend on the value $\mu(t)$ of a public information variable, which is drawn uniformly from the integers $1, \ldots, P$. In other words, agents have *trading strategies* which prescribes to agent i a bid $a_i^\mu \pm 1$ for each information μ. Each agent is assigned one such strategy, randomly chosen from the set of 2^P possible strategies of this type. Agents are adaptive and may decide to refrain from playing if their strategy is not good enough. More precisely, the bids of agents take the form $b_i(t) = \phi_i(t) a_i^{\mu(t)}$ where $\phi_i(t) = 1$ or 0 according to whether agent i trades or not. In order to assess the performance of their strategy, agents assign scores $U_i(t)$ which they update by

$$U_i(t+1) = U_i(t) - a_i^{\mu(t)} A(t) - \epsilon_i. \tag{71}$$

where

$$A(t) = \sum_{i=1}^{N} b_i(t) = \sum_{i=1}^{N} \phi_i(t) a_i^{\mu(t)}. \tag{72}$$

So if $-a_i^{\mu(t)} A(t)$ is large enough, i.e., larger than ϵ_i, the score U_i increases. The larger U_i, the more likely it is that the agent trades ($\phi_i = 1$). Here we suppose that

$$\mathrm{Prob}\{\phi_i(t) = 1\} = \frac{1}{1 + e^{\Gamma U_i(t)}} \tag{73}$$

where $\Gamma > 0$ is a constant. A good strategy prescribes bids a_i^μ which tend to coincide with those $b(t) = -\mathrm{sign}\, A(t)$ of the minority of agents. The connection with markets is realized assuming that $A(t)$ is proportional to the difference of price logarithms, i.e. $\log p(t+1) = \log p(t) + A(t)$.

The threshold ϵ_i in Eq. (71) models the incentives of agents for trading in the market. Some investors may have incentives to trade because they need the market for exchanging goods or assets. This corresponds to $\epsilon_i < 0$. On the contrary, speculators who only trade for profiting of price fluctuations typically have $\epsilon_i > 0$. Of course there may be prudent investors with $\epsilon_i > 0$ or risk-lover speculators with $\epsilon_i < 0$ and a whole range of other type of traders. Here we focus, for simplicity, on the case

$$\begin{aligned} \epsilon_i &= \epsilon \qquad \text{for } i \le N_s \\ \epsilon_i &= -\infty \qquad \text{for } N_s < i \le N \end{aligned}$$

The $N_p = N - N_s$ agents who have $\epsilon_i = -\infty$ — we call them *producers* after Refs. [Challet et al., 2000d, Zhang, 1999] — trade no matter what, whereas the remaining N_s — the *speculators* — trade only if the cumulated performance of their active strategy increases more rapidly than ϵt.

The collective behavior of the model in the stationary state is best understood in terms of the market's predictability, of the volatility and the number of active traders.

If the conditional time average $\langle A|\mu\rangle$ of $A(t)$ given $\mu(t) = \mu$ is non-zero, then the knowledge of $\mu(t)$ allows a statistical prediction of the sign of $A(t)$. A measure of predictability is hence given by

$$H_0 = \overline{\langle A\rangle^2} = \frac{1}{P}\sum_{\mu=1}^{P}\langle A|\mu\rangle^2$$

where we introduced the notation $\overline{(\ldots)}$ for averages over μ ($\langle\ldots\rangle$ denotes averages on the stationary state). When $H_0 = 0$ the market is unpredictable or *informationally efficient*. Volatility is instead defined as $\sigma^2 = \overline{\langle A^2\rangle}$ and it measures market's fluctuations. A further quantity of interest is the number $N_{\rm act}(t) = \sum_i\langle\phi_i(t)\rangle$ of active speculators in the market.

Exact results can be obtained in the thermodynamic limit, which is defined as the limit $N_s, N_p, P \to \infty$, keeping constant the reduced number of speculators and producers $n_s = N_s/P$, $n_p = N_p/P$. In this limit, both σ^2 and H_0 diverge with the system size, since $A(t) \sim \sqrt{N}$. Hence we shall consider the rescaled quantities H_0/P or σ^2/P. The calculation follows that of the standard MG. Here we just discuss the main step and the results. Following Ref. [Marsili and Challet, 2001a], we derive an Ito stochastic differential equations for the strategy scores $y_i(\tau) = U_i(t)$ in the rescaled continuous time $\tau = t/N$

$$\frac{{\rm d}y_i}{{\rm d}\tau} = -\overline{a_i\langle A\rangle_y} - \epsilon + \eta_i. \tag{74}$$

Here η_i is a zero average Gaussian noise term with

$$\langle\eta_i(\tau)\eta_j(\tau')\rangle = \frac{1}{N}\overline{a_ia_j\langle A^2\rangle_y}\delta(\tau-\tau'). \tag{75}$$

In Eqs. (74,75) averages $\langle\ldots\rangle_y$ are taken on the distribution of $\phi_i(t)$ in Eq. (73), which depends on $y_i(\tau)$ in a non-linear way: $\text{Prob}\{\phi_i(t) = 1\} = 1/[1 + e^{\Gamma y_i(\tau)}]$. Hence Eq. (74) is a quite complex system of non-linear equations with a noise strength proportional to the time dependent volatility $\overline{\langle A^2\rangle_y}$. This feedback will be responsible for the emergence of volatility build-ups.

Following Refs. [Challet et al., 2000c, Marsili and Challet, 2001a] we find that the fraction $\langle\phi_i\rangle$ of times that agent i plays his active strategy in the stationary state is the solution of the minimization of the function

$$H_\epsilon = \frac{1}{P}\sum_{\mu=1}^{P}\left[\sum_{i=1}^{N}\langle\phi_i\rangle a_i^\mu + \sum_{i=N_s+1}^{N_s+N_p} a_i^\mu\right]^2 + 2\epsilon\sum_i\langle\phi_i\rangle \tag{76}$$

with respect to $\langle\phi_i\rangle$. Note that for $\epsilon = 0$ this function reduces to the predictability H_0. For $\epsilon \neq 0$, the solution to this problem, and hence the stationary state, is unique. An exact statistical mechanics description of the solution $\{\langle\phi_i\rangle\}$ can be carried out with the replica method, because the replica symmetric ansatz is exact. Furthermore the solution to the Fokker-Planck equation corresponding to Eq. (74) can be well approximated by a factorized ansatz for $\epsilon \neq 0$. This means that the off-diagonal correlations vanish $[\langle(\phi_i - \langle\phi_i\rangle)(\phi_j - \langle\phi_j\rangle)\rangle = 0$ for $i \neq j]$ and, as a consequence, the volatility turns out to be given by $\sigma^2 = \overline{\langle A^2\rangle} = H_0 + \sum_{i=1}^{N_s}\langle\phi_i\rangle(1-\langle\phi_i\rangle)$. The solution $\{\langle\phi_i\rangle\}$ of the minimization

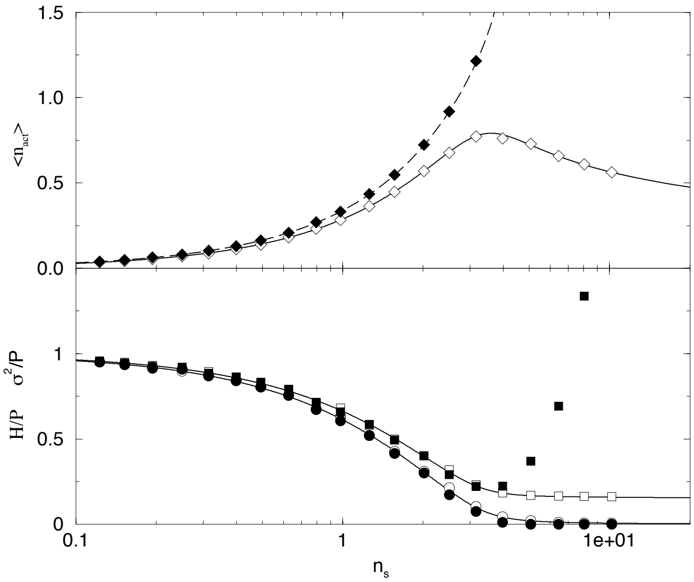

Figure 11: Theory and numerical simulations: n_{act} (top) and σ^2/P and H/P (bottom) as a function of n_s for $\epsilon = 0.1$ (solid line) and $\epsilon = -0.01$ (dashed line). Numerical results for $\epsilon = 0.1$ (open symbols) and $\epsilon = -0.01$ (full symbols) are averages over 200 runs, with $N_sP = 10000$ fixed and $\Gamma = \infty$.

of H_ϵ provides a complete description of the model in the limit $N \to \infty$ for $\epsilon > 0$. In particular the behavior is independent of Γ.

Fig. 11 shows that all these conclusions are perfectly supported by numerical simulations: With a fixed number n_p of producers, as the number n_s of speculators increases, the market becomes more and more unpredictable, i.e. H_0 decreases. At the same time also the volatility σ^2 decreases. In a market with few speculators ($n_s < 1$ in Fig. 11), most of the fluctuations in $A(t)$ are due to the random choice of $\mu(t)$ (i.e. $\sigma^2 \simeq H_0$) and the number n_{act} of active speculators grows approximately linearly with n_s.

When n_s increases further, the market reaches a point where it is barely predictable. Then, for $\epsilon > 0$ the number of active traders decreases and finally converges to a constant. This means that the market becomes highly selective: Only a negligible fraction of speculators trade ($\phi_i(t) = 1$) whereas the majority is inactive ($\phi_i(t) = 0$). The volatility σ^2 also remains constant in this limit.

For $\epsilon < 0$ we see a markedly different behavior: The number of active speculators continues growing with n_s even if the market is unpredictable $H_0 \approx 0$. The volatility σ^2 has a minimum and then it increases with n_s in a way which depends on Γ. In other words, $\epsilon = 0$ for $n_s \geq n_s^\star(n_p)$ ($= 4.15\ldots$ for $n_p = 1$) is the locus of a first order phase transition across which N_{act} and σ^2 exhibit a discontinuity. This same picture applies to a wider range of GCMG models such as that of Ref. [Challet et al., 2001].

Numerical simulations reproduce anomalous fluctuations similar to those of real financial markets close to the phase transition line. As shown in Fig. 12, the distribution of $A(t)$ is Gaussian for small enough n_s, and has fatter and fatter

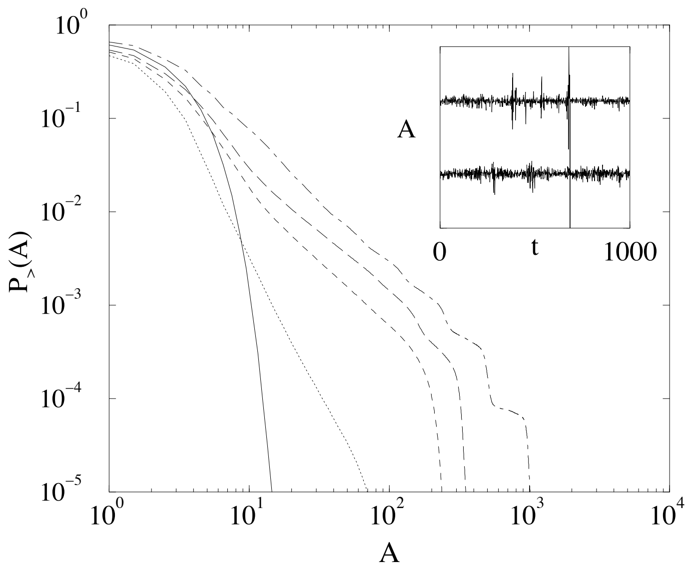

Figure 12: Probability distribution of $A(t) > 0$ for $n_s = 10$ (continuous line), 20,50,100,200 (dash-dotted line) ($PN_s = 16000$, $n_p = 1$, $\epsilon = 0.01$, $\Gamma = \infty$). Inset: time series of returns $A(t)$ showing volatility clustering for $n_s = 20$ (lower curve), but not for $n_s = 200$ (upper curve).

tails as n_s increases; the same behavior is seen for decreasing ϵ. In particular the distribution of $A(t)$ shows a power law behavior $P(|A| > x) \sim x^{-\beta}$ with an exponent which we estimated as $\beta \simeq 2.8, 1.4$ for $n_s = 20, 200$ respectively and $\epsilon = 0.01$. Note that a realistic value $\beta \approx 3$ is obtained for $n_s = 20$.

This is inconsistent, at first sight, with the theoretical results discussed previously for $N \to \infty$. Indeed, if the distribution of ϕ_i factorizes, $A(t)$ is the sum of N_s independent contributions and it satisfies the Central Limit Theorem. This implies that for $\epsilon \neq 0$ the variable $A(t)/\sqrt{N}$ converges in distribution to a Gaussian variable with zero average and variance σ^2/N in the limit $N \to \infty$. There are no anomalous fluctuations and no stylized facts. Fig. 13 indeed shows that the anomalous fluctuations of Fig. 12 are finite size effects which disappear abruptly as the system size increases (or if Γ is small).

In order to understand these finite size effects, we note that volatility clustering arises because the noise strength in Eqs. (74,75) is proportional to the time dependent volatility $\overline{\langle A^2 \rangle_y}$. The noise term is a source of correlated fluctuations because $\overline{a_i a_j \langle A^2 \rangle_y}/N \sim 1/\sqrt{N}$ is small but non zero, for $i \neq j$. It is reasonable to assume that the dynamics will sustain collective correlated fluctuations in the y_i only if the correlated noise is larger than the signal $-\overline{a_i \langle A \rangle_y} - \epsilon$ which agents receive form the deterministic part of Eq. (74). Time dependent volatility fluctuations would be dissipated by the deterministic dynamics otherwise. A quantitative translation of this insight goes as follows: The noise correlation term is of order $\overline{a_i a_j \langle A^2 \rangle_y}/N \sim \sigma^2/P^{3/2}$ for $i \neq j$. This should be compared to the square of the deterministic term of Eq. (74) $[\overline{a_i \langle A \rangle_y} + \epsilon]^2 \sim \left[\sqrt{H_0/P} + \epsilon\right]^2$.

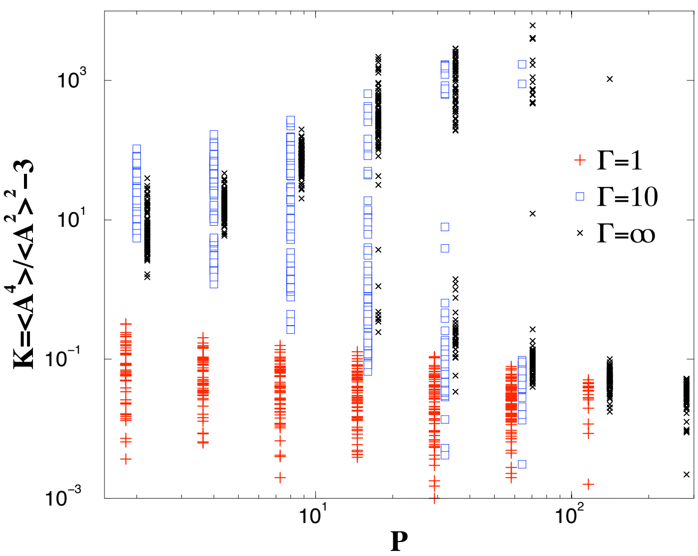

Figure 13: Kurtosis of $A(t)$ in simulations with $\epsilon = 0.01$, $n_s = 70$, $n_p = 1$ and several different system sizes P for $\Gamma = 1, 10$ and ∞.

Rearranging terms, we find that volatility clustering sets in when

$$\frac{H_0}{\sigma^2} + 2\epsilon\sqrt{\frac{H_0}{P}}\frac{P}{\sigma^2} + \epsilon^2\frac{P}{\sigma^2} \simeq \frac{K}{\sqrt{P}} \tag{77}$$

where K is a constant. This prediction is remarkably well confirmed by Fig. 14: In the lower panel we plot the two sides of Eq. (77) as a function of n_s, for different system sizes. The upper panel shows that the volatility σ^2/N starts deviating from the analytic result exactly at the crossing point $n_s^c(P)$ where Eq. (77) holds true. Furthermore the inset shows that the region $n_s > n_s^c(P)$ is described by a novel type of scaling limit. Indeed the curves of Fig. 14 collapse one on top of the other when plotted against $n_s/n_s^c(P)$.

The non-linearity of the response of agents is crucial for the onset of volatility time dependence. If Γ is small the response becomes smooth and anomalous fluctuations disappear (see Fig. 13).

The fact that, in finite systems, stylized facts arise only close to the phase transition is reminiscent of finite size scaling in the theory of critical phenomena: In d-dimensional Ising model, for example, at temperature $T = T_c + \varepsilon$ critical fluctuations (e.g. in the magnetization) occur as long as the system size N is smaller than the correlation volume $\sim \varepsilon^{-d\nu}$. But for $N \gg \varepsilon^{-d\nu}$ the system shows the normal fluctuations of a paramagnet.

Eq. (77) and $H_0/P \sim \epsilon^2$ imply that the same occurs in the GCMG with $d\nu = 4$. In other words, the critical window shrinks as $N^{-1/4}$ when $N \to \infty$. However, because of the long range nature of the interaction, anomalous fluctuations either concern the whole system or do not affect it at all, as clearly shown in Fig. 13. In the critical region the Gaussian phase coexists probabilistically with a phase characterized by anomalous fluctuations. This and the discontinuous nature of the transition at $\epsilon = 0$, are usually typical of first order phase transitions.

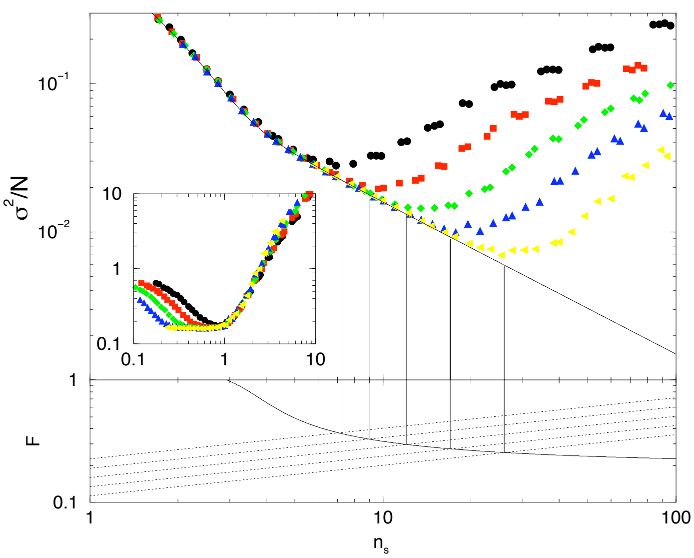

Figure 14: Onset of the anomalous dynamics for different system sizes. Top: σ^2/N for different series of simulations with $L \equiv PN_s$ constant: $PN_s = 1000$ (circles), 2000 (squares), 4000 (diamonds), 8000 (up triangles) and 16000 (left triangles). In all simulations $n_p = 1$, $\epsilon = 0.1$ and $\Gamma = \infty$. Bottom: L.H.S. of Eq. (77) (full line) from the exact solution and $K/\sqrt{P} = K(n_s/L)^{1/4}$ (parallel dashed lines) as a function of n_s ($K \simeq 1.1132$ in this plot). The intersection defines $n_s^c(P)$. Inset: Collapse plot of σ^2/N as a function of $n_s/n_s^c(P)$.

The picture of a phase transition controlled by the signal to noise ratio appears to be universal for Minority Games. Finite size effects close to the phase transition of the standard MG [Challet and Zhang, 1997] are indeed explained by the same generic argument: When the signal to noise ratio H_0/σ^2 is of order $1/\sqrt{P}$ self-sustained collective fluctuations arise.

Volatility clustering in real markets is known to be due to wild fluctuations in the volume of trades. Volume is the number of active traders $N_{\text{act}} + N_p$ in the GCMG. Hence wild volume fluctuations require correlated collective fluctuations in the behavior of agents which only arise close to criticality. This suggests that real markets operate close to a phase transition. Numerical simulations suggest that exponents vary continuously on the line of critical points. This raises the question of why real markets self-organize close to the critical surface with $\alpha \approx 3$.

We conclude that the GCMG exhibits a quite peculiar type of phase transition which mixes properties of continuous and discontinuous transitions. Finite size effects clearly relate the occurrence of stylized facts to the analytic nature of the phase transition. The extension of renormalization group approaches to this system promises to be a quite interesting challenge.

9 Improving cooperation in the Minority Game

The Minority Game addresses the question of how efficiently a group of adaptive agents competing for limited resources can coordinate. The fact that agents with limited rationality and resources can perform better than by taking random decisions was regarded as a non-trivial result in early papers. This led to a 'quest for efficiency' in a series of papers aiming to find learning dynamics leading to better coordination (see e.g. [Hod and Nakar, 2002, Reents et al., 2001, Paczuski et al., 2000, Kinzel et al., 2000, Wakeling and Bak, 2001]).

All these works, however, do not directly pose the question of whether the agents, given their constrained capacities, play optimally or not. In game theoretic terms, playing optimally means that agents should be in a Nash equilibrium—a state where no agent can improve his payoff by adjusting his strategy if others stick to their choices. It turns out that the agents in the standard Minority Game do not converge to a Nash equilibrium because they do not play *strategically.* Putting it simply, agents behave as if they were playing against a *market* and not against $N-1$ other agents. The key issue is that this behaviour neglects the impact that each player has on the market. A minimal requirement for a strategic behavior is to consider the effects of one's own actions on oneself. This turns out to be enough to 'correct' the learning dynamics that then converges to a Nash equilibrium [Challet et al., 2000c, Marsili et al., 2000]. In the dynamical equations of the resulting game, the market impact plays a role similar to the cavity, self-interaction term—also called after Onsager—in spin glasses [Mézard et al., 1987]. Removing this term causes replica symmetry breaking, leading to a totally novel scenario characterized by exponentially many (in N) possible equilibria [Martino and Marsili, 2001, Heimel and Martino, 2001, Marsili et al., 2001].

In the standard Minority Game, the global efficiency σ^2, which also measures the total agent loss, varies but remains proportional to N. Is it the best that the agents can do? Imagine a situation where N is odd and $(N+1)/2$ agents stick to action $a_i = +1$ whereas the rest take $a_i = -1$. This simple arrangement has

$\sigma^2 = 1$ which is way better than what we found previously. Why do the agents not consider this? How can the agents coordinate to such an arrangement[22]?

It is preferable to address these questions in simple Minority Game models first, avoiding the complications of dealing with strategies, information (whether endogenous or exogenous), hence disorder.

One such simple model was proposed by Reents *et al.* [Reents et al., 2001]. In their model winners of the last game stick to their choice, losers individually change their minds with probability p. In other words agents react to an individual piece of information, which is whether their choice was the right one or not.

This model is easy to understand: as there are $(N + |A|)/2$ losers and $|A|$ is at most of order N, the average number of changes is $p(N + |A|)/2 \ \sim pN$. Three regimes can be distinguished:

- $pN = x =$ constant: the number of people that change their mind does not depend on N, hence this leads to very small fluctuations $\sigma^2 \sim O(1)$.
- $pN \sim O(\sqrt{N})$: the fluctuations are typically of order N, which is what happens in the original Minority Game with histories.
- $p \sim O(1)$: the fluctuations are of order N^2. $A(t)$ is characterized by a double peaked distribution.

Since the model is very simple, it can be tackled analytically. We refer to Ref. [Reents et al., 2001] for details.

Refs. [Marsili and Challet, 2001b] observes that there are two separated issues in the Minority Game. One is the competitive aspect by which the agents try to exploit asymmetries in the game's outcome $A(t)$. This is the force driving to information efficiency $\langle A \rangle \approx 0$ and it has to do with predictability. The other is the coordination aspect of the game and it is related to the volatility σ^2. Loosely speaking, no agent in the Minority Game likes volatility. Furthermore volatility increases the fluctuations in the behaviour of agents and this feeds back into the collective behaviour causing volatility build-ups.

The analysis of Ref. [Marsili and Challet, 2001b] was later refined in paper [Marsili, 2001] (see section 5). It shows that very simple models are able to explain puzzling features of the Minority Game in its full-blown complexity, such as the dependence of aggregate fluctuations (σ^2) on microscopic randomness (Γ) and on initial conditions.

In particular, in the extreme case when the spread of initial conditions $\Delta_i(0)$ is very large[23] σ^2 becomes of order 1, in agreement with the numerical simulations of Refs. [Garrahan et al., 2000, Marsili and Challet, 2001a] for $\alpha \to 0$. This means that agents, starting from extremely different initial conditions, finally split into two groups of equal size playing opposite actions. This outcome, however, does not come by the virtue of agents' ability to coordinate but is rather already buried in the initial conditions: roughly speaking, half of the population is convinced, right from the beginning, that $a_i = +1$ is way better than $a_i = -1$ whereas the others have opposite beliefs. Agents who learn to

[22]Note that there are very many such arrangements.

[23]A more precise condition is that the number of agents with initial conditions in any interval of size $\delta\Delta$ be finite, as $N \to \infty$.

coordinate should form these beliefs endogenously in the course of the game[24]. But why do the agents in the Minority Game fail to do this?

9.1 Market impact and Nash equilibria

As explained above, the optimal way in which agents can play the minority game is to split into two equally sized groups taking opposite actions. If N is odd $(N+1)/2$ agents, being in the majority, lose. But each of them cannot do better by changing side: the majority side would move with heras sheswitches side (unless someone else change her mind as well). No one can improve her situation by unilaterally deviating from her behaviour. This situation is exactly what game theory calls a Nash equilibrium [Fudenberg and Tirole, 1991].

The Nash equilibria of the simple minority game discussed earlier, where strategies are just actions $a_i = \pm 1$, were first discussed in Ref. [Marsili and Challet, 2001a]. Briefly, all arrangements where $N - 2k$ agents play *mixed strategies*, i.e. choose $a_i = \pm 1$ at random, and the remaining $2k$ agents split into two groups of equal size taking opposite actions, are Nash equilibria. These Nash equilibria have $\sigma^2 = N - 2k$. Those k agents taking the action $a_i = +1$ can be choosen in $\binom{2k}{k}$ number of ways out of the $2k$. The number of Nash equilibria grows as 2^N; it becomes huge already for moderately large N. Those Nash equilibria with largest $k \leq N/2$ are the most efficient, with $\sigma^2 \leq 1$.

Why do agents fail to reach these optimal states? The reason is that *agents in all versions of the Minority Game discussed so far do not play a game against $N-1$ other agents. Rather, they behave as if they were playing against the process $A(t)$*. The problem becomes evident for $N = 1$: in this extreme Minority Game the agent would continue endlessly to react to herself, switching side at each round. A strategic player would do no better in terms of payoff, as shewould lose whatever side shechooses. But at least shewould realize that there is no way out. It may seem strange at first sight, but the fact that agents react to themselves is what hinders them from reaching an optimal Nash equilibrium even for large N. Even though for large N each agent's contribution to $A(t)$ is small, if all agents neglect it, they will fail to reach the Nash equilibrium.

To better understand this issue let us go back to the discussion of section 5 on the Minority Game without information. There agents learning through Eq. (27) behave as price takers: They totally neglect the fact that price changes – i.e. $A(t)$ – also depend on their choice $a_i(t)$. This may seem reasonable given that agents are very many and the impact of each of them is very small. As we shall see in a while (see also Refs. [Challet et al., 2000c, Marsili et al., 2000, Marsili and Challet, 2001a]), this argument is misleading because indeed price taking behavior has very strong consequences. Let us consider a slightly different learning dynamics:

$$\Delta_i(t+1) = \Delta_i(t) - \frac{\Gamma}{N}[A(t) - \eta a_i(t)]. \tag{78}$$

The η term in Eq. (78) describes the fact that agent i accounts for his own contribution to $A(t)$. For $\eta = 1$ indeed, agent i considers only the behavior of other agents $A(t) - a_i(t)$ and does not react to his own action $a_i(t)$. In other words, η measures the extent to which agents account for their "market impact".

[24] This seems to be the intuition of Arthur when he states that "Expectations will be forced to differ" [Arthur, 1994] in contexts such as the *El Farol* bar problem.

For $\eta = 0$ we recover the results discussed in section 5. But the situation changes drastically as soon as agents start to account for their market impact, i.e. for $\eta > 0$. To see this, let us take the average of Eq. (78) in the long time limit and define $m_i = \langle a_i \rangle$. We note that

$$\langle \Delta_i(t+1) \rangle - \langle \Delta_i(t) \rangle = -\frac{\Gamma}{N} \left[\sum_{j \in N} m_j - \eta m_i \right] = -\frac{\Gamma}{N} \frac{\partial H_\eta}{\partial \eta} \tag{79}$$

where

$$H_\eta = \frac{1}{2} \left(\sum_{i \in N} m_i \right)^2 - \frac{\eta}{2} \sum_{i \in N} m_i^2. \tag{80}$$

A close inspection[25] of these equations implies that m_i are given by the minima of H_η.

Note that H_1 is an Harmonic function of m_i's. Hence it attains its minima (and maxima) on the boundary of the hypercube $[-1, 1]^N$. So for $\eta = 1$ all agents take always the same actions $a_i(t) = m_i = +1$ or -1 and the waste of resources is as small as possible: Indeed $\sigma^2 = 0$ or 1 if N is even or odd, which is a tremendous improvement with respect to the case $\eta = 0$ (where $\sigma^2 \sim N$ or N^2). These states are indeed Nash equilibria [Marsili and Challet, 2001b] of the associated N persons minority game. This argument extends to all $\eta > 0$: The stationary states of the learning process for any $\eta > 0$ are Nash equilibria[26]. Hence as soon as agents start to account for their market impact ($\eta > 0$) the collective property of the system changes abruptly and inefficiencies σ^2 are drastically reduced.

Again the asymptotic state is not unique and it is selected by the initial conditions. However now the set of equilibria is discrete and the system jumps discontinuously from an equilibrium to another, as the initial conditions $\Delta_i(0)$ vary. This contrasts with the $\eta = 0$ case, where the equilibrium shifts continuously as a function of the initial conditions [Marsili and Challet, 2001a].

The issue of market impact in Minority Game was first raised in paper [Challet et al., 2000c] and in Ref. [Marsili et al., 2000] for the model in its full complexity. There Minority Game agents were called *naïve* as opposed to the *sophisticated* strategic players of game theory. Agents in the Minority Game naïvely neglect their market impact, assuming that it is negligible. Considering the Minority Game as a market model, they behave as *price takers*. This assumption about traders may be realistic: traders may really behave that way. But the results show that the assumption is by no means an innocent one. If all agents account for their impact the collective behaviour changes dramatically.

In particular, if agents correctly account for their market impact, they reach a Nash equilibrium.

---

[25] The first order conditions on H_η imply that if $-1 < m_i < 1$ then $\langle \Delta_i(t+1) \rangle = \langle \Delta_i(t) \rangle$, i.e. the process $\Delta_i(t)$ is stationary. Else if $m_i = +1$ (or -1) one should have $\Delta_i(t) \to +\infty$ (or $-\infty$), which is precisely what constrained minimization and Eq. (79) say.

[26] The proof of this statement goes as follows: Suppose that m^* is an equilibrium with $-1 < m_k^* < 1$ for $k = 1, \ldots, n$ and $m_i^* = \pm 1$ for $i > n$. The conditions for a minimum requires that H_η is locally positive definite around m^*. At least n eigenvalues of the matrix $\frac{\partial^2 H_\eta}{\partial m_i \partial m_j}$ must be non-negative. But this matrix has only one positive eigenvalue $\lambda = N - \eta$ and $N - 1$ negative eigenvalues $\lambda = -\eta$. Hence n can at most be 1, which can occur for N odd.

9.2 Nash equilibria of the Minority Game

A Nash equilibrium is defined in terms of the strategies s_i which players can choose and in terms of the payoff matrix $u_i(s_i, s_{-i})$, where $s_{-i} = \{s_j,\ j \neq i\}$ is the usual game theoretic notation for the strategies of opponents. In the Minority Game the payoff matrix is given by

$$u_i(s_i, s_{-i}) = -\frac{1}{N}\sum_{j=1}^{N} \overline{a_{i,s_i} a_{j,s_j}} = -\frac{1}{NP}\sum_{\mu=1}^{P}\sum_{j=1}^{N} a^{\mu}_{i,s_i} a^{\mu}_{j,s_j} \tag{81}$$

where $a^{\mu}_{i,s}$ are the randomly drawn look-up tables of agents.

The game theoretic interpretation of the Minority Game is discussed in Ref. [Marsili et al., 2000]. In brief, we imagine we deal with a single stage game with a state of the world μ which is drawn at random and Eq. (81) is the expected payoff[27]. Eq. (81) is also the expected payoff of a game where each player i is randomly matched with another one (j) to play a game with payoffs $u_i^{\mu}(s_i, s_j) = -a^{\mu}_{s_i,i} a^{\mu}_{s_j,j}$ where the state $\mu = 1, \ldots, P$ of Nature is drawn also randomly. Finally, game payoffs similar to those of Eq. (81) are also found in contexts where N agents compete for the exploitation of P exhaustible resources. Urban traffic, as shown in [Martino et al., 2003], is one example.

Independently of its interpretation, the payoff matrix in Eq. 81 represents an interesting instance of a complex system of interacting heterogeneous agents whose rich behaviour deserves investigation in its own right. Ref. [Marsili et al., 2000] found that Nash equilibria in evolutionarily stable strategies are the minima of σ^2 and each player plays pure strategies. This implies $H = \sigma^2$ (see e.g. Eq. 39). The study of Nash equilibria was refined in Ref. [Martino and Marsili, 2001]: it turns out that the characterization of the optimal Nash equilibrium, that with minimal σ^2, requires full replica symmetry breaking in the statistical mechanics approach. In simple terms this signals the existence of a complex hierarchical organization of the Nash equilibria. Ref. [Martino and Marsili, 2001] shows that the number of Nash equilibria grows exponentially with N, with a growth rate Σ which depends on α as shown in Fig. 9.2.

The properties of the Nash equilibria (NE) are quite different from the stationary states of the Minority Game (with naïve agents). In particular:

- Nash equilibria are local minima of σ^2 whereas the stationary state of the Minority Game with naïve agents corresponds to minimum of H.
- There are (exponentially) many disconnected Nash equilibria. Naïve agents reach a single stationary state, degenerate on a connected set for $\alpha < \alpha_c$.

[27] In the Minority Game as well as in the El Farol bar problem, one imagines that the interaction is repeated in time, neglecting all the strategic intricacies which arise from the inter-temporal nature of repeated games (such as reputation, punishment or signalling). It is reasonable to assume that complications of this sort play a marginal role in a context such as the one we are interested in, where the number of agents is very large. Actually El Farol bar goers consider it just too complex to undertake a strategic behaviour even in each single stage. Each of them could be better off deciding her action a_i at each stage. If they resort to the recommendations $a^{\mu}_{s,i}$ of their look-up tables, it means that there are implicit computational and implementation costs which make them prefer simple, though sub-optimal, rules of thumb. This is more or less the same attitude as agents in the original Minority GameIn principle, players may be forced to one of the types of behaviour prescribed by their repertoire of look-up tables by some other constraints. In this case, it is important that players do not know the value of μ before deciding which strategy to use, because otherwise they could take a decision s_i^{μ} which depends on μ.

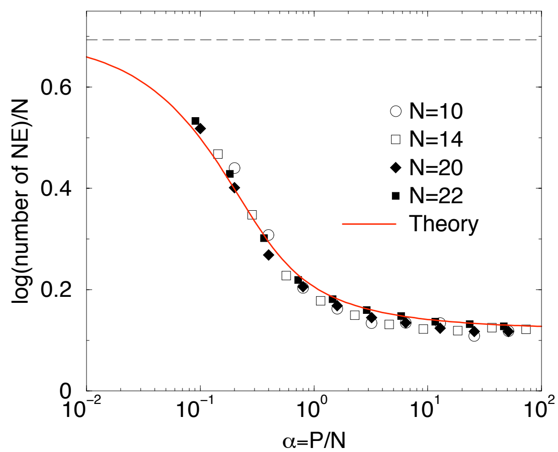

Figure 15: Logarithm of the average number of Nash Equilibria divided by N as a function of α.

- Sophisticated players use pure strategies in NE whereas naïve agents alternate strategies. As a result $H = \sigma^2$ in NE whereas $H < \sigma^2$ in the original Minority Game.

- The NE to which the learning dynamics of sophisticated agents converges is selected by the initial conditions. As the initial conditions vary the NE changes discontinuously. On the other hand, the learning dynamics of naïve agents converges to a stationary state which for $\alpha < \alpha_c$ is unequivocal and depends *continuously* on the initial conditions.

- In the stationary state of the Minority Game with naïve agents for $\alpha < \alpha_c$, a feedback of fluctuations from microscopic degrees of freedom to macroscopic quantities leads to a dependence of σ^2 on the learning rate Γ. This effect is absent for all values of α from the learning dynamics of sophisticated agents, who converge to a NE with no fluctuations.

- The behaviour of the Minority Game with naïve agents is qualitatively the same both under endogenous and exogenous information (see paper [Cavagna, 1999] but also Ref. [Challet and Marsili, 2000]). In contrast, the learning dynamics of sophisticated agents under endogenous information is very different from that under exogenous information[28]. Hence the origin of information is not irrelevant when agents account for their market impact.

The fact that sophisticated agents have so many Nash equilibria at their disposal where to converge to raises interesting questions: how do agents manage to coordinate and select the same Nash equilibrium? How much time do agents need to "learn" it? These questions are crucial in realistic cases where agents

[28] In few words—we refer to Ref. [Marsili et al., 2000] for more details—the dynamics of $\mu(t)$ become deterministic because all agents get frozen, and it locks into a periodic orbit of period $\sim \sqrt{P}$. Almost all information patterns μ', those not on the orbit, are never visited in the stationary state.

have a finite memory, i.e. forget about the past. Having a finite memory may seem a further limitation but is actually a necessity to adapt optimally if the environment is changing (either because interactions change or because other agents change). A finite memory is modeled introducing a forgetting rate λ into the dynamics,

$$U_{i,s}(t+1) = \left(1 - \frac{\lambda}{N}\right) U_{i,s}(t) + \frac{\lambda}{N} u_i[s, s_{-i}(t)] \tag{82}$$

by which agents learn about the payoffs $u_i[s, s_{-i}(t)]$ that they receive from strategy s. Ref. [Marsili et al., 2001] has found a phase transition from a phase where agents manage to coordinate efficiently, when their memory extends sufficiently far into the past ($\lambda < 0.46\Gamma$) to a random phase where no coordination takes place. This transition, which is continuous in a stationary setting, becomes discontinuous in a "changing world". As Ref. [Marsili et al., 2001] puts it, the occurrence of a dynamical transition "is further evidence that an analysis in terms of Nash equilibria may not be enough to predict the collective behaviour of a system. Agents may fail to coordinate [to] Nash equilibria because of purely dynamical effects."

9.3 From naïve to sophisticated agents

That players should stick to just one strategy—the best one—is intuitive at first sight. Indeed the nature of the game is very similar to that of typical coordination games [Bottazzi et al., 2003]. If a player has an optimal strategy, adopting this strategy is the best that he can do to reduce σ^2. This is the choice which the opponents welcome the most and that they anticipate. There is no reason why the agents should resort to mixed strategies in the Minority Game.[29] The reason why naïve agents do not stick to a single strategy in the usual Minority Game has nothing to do with game theory. Rather they do so because they neglect their market impact.

In order to see this, and to understand how to relate the Minority Game behaviour with that of Nash equilibria—which appears completely different—paper [Challet et al., 2000c] suggested studying a learning dynamics

$$U_{i,s}(t+1) = U_{i,s}(t) - a_{i,s}^{\mu(t)} A(t) + \eta \delta_{s,s_i(t)} \tag{83}$$

with an additional term rewarding the strategy actually played by agent i. With $\eta = 0$ this clearly reproduces the behaviour of naïve agents in the Minority Game, whereas with $\eta = 1$ it reproduces the behaviour of sophisticated agents[30].

Fig. 16 shows the dramatic influence that η has one the fluctuations in the stationary state. In the information rich phase ($\alpha > \alpha_c$) global efficiency improves smoothly with the reward η. This corresponds to crossing the full line boundary between the two phases in Fig. 17. But in the informational efficient phase ($\alpha < \alpha_c$), even an infinitesimal reward is able to reduce volatility by a

[29]The typical case where agents have strict incentives to randomize, i.e. to play a mixed strategy, is the matching penny game where disclosing information about your choice makes your opponent win, whereas you lose (e.g. goal-keeper/penalty shooter interaction).

[30]To be precise, the last term should be $+\eta a_{i,s}^{\mu(t)} a_{i,s_i(t)}^{\mu(t)}$, as observed in Mansilla [Mansilla, 2000]. However, this has the same effect as $+\eta\delta_{s,s_i(t)}$ on the long term behaviour, because $\overline{a_{i,s} a_{i,s_i(t)}} \cong \delta_{s,s_i(t)}$.

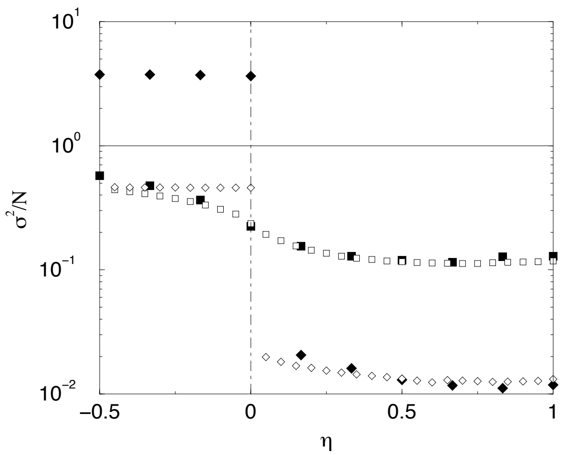

Figure 16: Effect of accounting for market impact in the Minority Game with $\Gamma = \infty$ (full symbols) and $\Gamma \ll 1$ (open symbols). Squares refer to $\alpha > \alpha_c$ whereas diamonds to $\alpha < \alpha_c$.

finite amount (which amounts to crossing the dashed line in Fig. 17). The effect is more spectacular if agents are very reactive ($\Gamma = \infty$) than if they learn at a low rate ($\Gamma \ll 1$). What is remarkable is that the phase transition at α_c completely changes its nature for $\eta > 0$ and it disappears when agents account properly for the market impact ($\eta = 1$). This is what we alluded to when we stated, at the end of the previous chapter, that stylized facts disappear if agents account for their market impact.

Understanding what happens in the (α, η) space of this generalized Minority Game model requires some technicalities which we deal with in the next subsection, for the interested reader. The result is shown in Fig. 17 and can be summarized as follows: the properties of the Minority Game with naïve agents generalize to the whole RS region of the phase diagram. There we find a single (replica symmetric) stationary state, independent of initial conditions. Beyond the transition line in Fig. 17, the system has properties similar to those of Nash equilibria (replica symmetry breaking): there are many disconnected states which are selected by the choice of initial conditions. It is possible to show [Marsili et al., 2000] that σ^2 decreases and all individual payoffs increase when η increases. The transition is continuous for $\alpha > \alpha_c$, but it becomes discontinuous for $\alpha < \alpha_c$. The segment between the origin and the point $(\alpha_c, 0)$ is rather peculiar: the stationary state is degenerate on a continuous set and its properties change abruptly as this line is crossed. As anticipated, the interested reader can find more details in the next section.

9.4 The AT-MO line

The asymptotic regime of the dynamics provided by Eq. (83) have been related, in Ref. [Martino and Marsili, 2001], to the local minima of

$$\mathcal{H}_\eta = (1-\eta)H + \eta\sigma^2. \tag{84}$$

This work has shown that a phase transition separates, in the (α, η) plane,

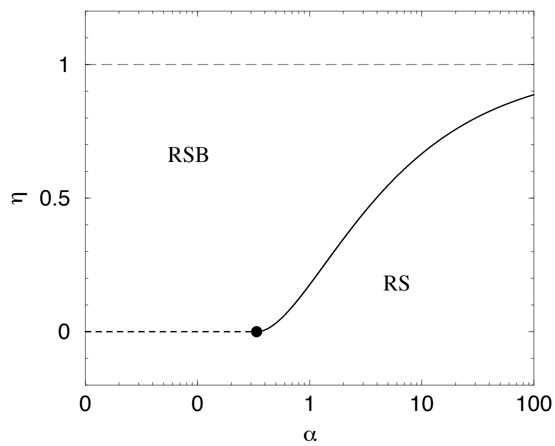

Figure 17: Phase diagram of the Minority Game in the α, η plane. The phase below the line is replica symmetric and the stationary state is unique whereas above the line there are many different stationary states. The transition is continuous across the full line ($\alpha > \alpha_c$) and discontinuous across the dashed thick line ($\alpha < \alpha_c$). The stationary state has a continuous degeneracy on the line $\eta = 0$ for $0 < \alpha < \alpha_c$. The light dashed line is the locus of Nash equilibria ($\eta = 1$).

the behaviour of the Minority Game ($\eta = 0$) from that of NE ($\eta = 1$). At this phase transition, replica symmetry breaks down, as shown by the calculation of the stability of the replica symmetric solution[31]. The line $\eta_c(\alpha)$ above which the replica symmetric solution become unstable is shown in Fig. 17 and it is usually called the *AT line* after de Almeida and Thouless [de Almeida and Thouless, 1978], who first discussed it for spin glasses.

In order to understand the meaning of this result, let us focus on the Minority Game with $S = 2$ strategies per agent. The replica symmetric solution is characterized by the equality $Q = q$ of the diagonal and off-diagonal overlaps. These are defined considering two replicas of the same system corresponding to two dynamical paths starting from different initial conditions. If m_i and m'_i are the averages of the strategic choices $s_i(t)$ in the two replicas, the overlaps are defined as

$$q = \frac{1}{N}\sum_{i=1}^{N} m_i m'_i, \qquad Q = \frac{1}{N}\sum_{i=1}^{N} m_i^2. \tag{85}$$

The equality $Q = q$ means that the two replicas are characterized by the same behaviour in the stationary state: $m_i = m'_i$. This indeed is what one should expect when the stationary state is unique. When the replica symmetric solution breaks down the off-diagonal overlap q takes, in principle, a whole range of values from $-Q$ to Q, signalling that beyond the line $\eta_c(\alpha)$ the minimum of $\mathcal{H}_\eta$, and hence the stationary state, are no longer unique. Different replicas, i.e. dynamical realizations starting from different initial conditions, may converge to different stationary states.

[31] Actually the transition line in Ref. [Martino and Marsili, 2001] is affected by a computational error, corrected in paper [Heimel and Martino, 2001].

The line $\eta_c(\alpha)$ can also be derived within the generating functional approach [Heimel and Coolen, 2001], as shown in paper [Heimel and Martino, 2001], where it is called the *memory onset* (MO) line. Indeed Heimel and De Martino show that above the MO line $\eta_c(\alpha)$ the dynamics acquire long term memory. In loose words, for $\eta < \eta_c(\alpha)$ the dynamics always "asymptotically forgets" perturbations in the early stages. On the contrary, for $\eta > \eta_c(\alpha)$, perturbation in the early stages of the dynamics may lead to a different stationary state: hence the system "remembers" early perturbations. This is exactly what one should expect if at $\eta_c(\alpha)$ the uniqueness of the minimum of $\mathcal{H}_\eta$ ceases to hold.

Indeed there is a very simple derivation of the AT or MO line. At $\eta_c(\alpha)$ the unique minimum of $\mathcal{H}_\eta$ must turn into a saddle point. $\mathcal{H}_\eta$ is a quadratic function of the m_i,

$$\mathcal{H}_\eta = c + \sum_i g_i m_i + \sum_{i,j} m_i T_{i,j} m_j \tag{86}$$

where c and g_i are constants and

$$T_{i,j} = \overline{\xi_i \xi_j} \left(1 - m_j^2\right) - \eta \overline{\xi_i^2} \left(1 - m_i^2\right) \delta_{i,j} \tag{87}$$

The minimum of $\mathcal{H}_\eta$ becomes a saddle point when the smallest of the eigenvalues of $\hat{T}$ becomes negative. Note that $\hat{T} = \hat{J} \cdot \hat{D}$ where $J_{i,j} = \overline{\xi_i \xi_j} - \eta \overline{\xi_i^2} \delta_{i,j}$ and $D_{i,j} = (1 - m_i^2)\delta_{i,j}$ is diagonal. In particular, $T_{i,j} = 0$ if j is a frozen agent ($m_j = \pm 1$), because $D_{j,j} = 0$. Hence each frozen agent corresponds to a trivial zero eigenvalue of the matrix $\hat{T}$. The non-trivial part of the eigenvalue spectrum of T is that relative to the sector of non-frozen—or *fickle*—agents. Let $N_u = (1 - \phi)N$ be the number of fickle agents. In order to focus on the sector of fickle agents, call $\hat{A}^{(u)}$ the $N_u \times N_u$ matrix obtained from a generic matrix $\hat{A}$ by deleting all the $N - N_u$ rows and columns corresponding to frozen agents. The smallest eigenvalue of $\hat{T}^{(u)}$ turns negative when the determinant of $\hat{T}^{(u)}$ vanishes. But $\det \hat{T}^{(u)} = \det \hat{J}^{(u)} \det \hat{D}^{(u)}$ and $\hat{D}^{(u)}$ is a positive definite matrix. The minimum of H_η becomes a saddle point when the smallest eigenvalue of $\hat{J}^{(u)}$ vanishes. The distribution of eigenvalues of matrices such as $\hat{J}^{(u)}$, which are proportional to $\overline{\xi_i \xi_j}$, is known (see e.g. [Sengupta and Mitra, 1999]). In particular, this allows us to compute the smallest eigenvalue of $\hat{J}^{(u)}$, which is given by

$$\lambda_- = \frac{1}{2} \left[\left(1 - \sqrt{\frac{1-\phi}{\alpha}}\right)^2 - \eta \right]. \tag{88}$$

Therefore λ_- vanishes when

$$\eta_c(\alpha) = \left(1 - \sqrt{(1-\phi)/\alpha}\right)^2 \tag{89}$$

which is exactly the AT or MO line[32].

[32] Paper [Heimel and Martino, 2001] reports the result $1 - \phi = \alpha \left[1 - \eta(1 + \chi)\right]^2$ which coincides with Eq. 89. Indeed, using the saddle point equations, one finds that $\chi = 1/\sqrt{\eta} - 1$ on the line $\eta_c(\alpha)$.

This calculation can also be read as the stability analysis of stationary states for the dynamics in the continuum time limit[33]. This shows that the replica symmetric solution becomes dynamically unstable for $\alpha < \alpha_c(\eta)$. Naively speaking, for each negative eigenvalue we have a bifurcation point of the dynamics which gives rise to the degeneracy of stationary states.

A Replica method for the MG

As in Ref. [Challet et al., 2000d], 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 $a^{\mu}_{+1,i}$, $a^{\mu}_{-1,i}$ for $i = 1,\ldots,N$ and $\mu = 1,\ldots,P$. 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:

$$\begin{aligned} P(a_+1, a_-1) &= \frac{c}{2}\left[\delta_{a_+1,+1}\delta_{a_-1,+1} + \delta_{a_+1,-1}\delta_{a_-1,-1}\right] \\ &+ \frac{1-c}{2}\left[\delta_{a_+1,-1}\delta_{a_-1,+1} + \delta_{a_+1,-1}\delta_{a_-1,+1}\right] \end{aligned} \tag{92}$$

   Note that, for $c = 0$ agents choose just one strategy a_+1 and fix $a_-1 = -a_+1$ as its opposite, whereas for $c = 1$ they have one and the same strategy $a_+1 = a_-1$. The original random case [Challet and Zhang, 1997, Savit et al., 1999] corresponds to $c = 1/2$. These agents assign scores $U_{s,i}(t)$ to each of their strategies and play the strategy $s_i(t)$ with the highest score, as discussed in the text. Therefore for speculators:

$$a_{\rm spec}(t) = a^{\mu(t)}_{s_i(t),i}. \tag{93}$$

2. then we consider $N^{\rm indep}_{\rm prod} = \rho N$ **producers**: They have only one randomly and independently drawn strategy b^{μ}_i so

$$a_{\rm prod}(t) = b^{\mu(t)}_i. \tag{94}$$

   Producers have a predictable behavior in the market and they are not adaptive. Instead of ρN *independent* producers one can also consider $N^{\rm dep}_{\rm prod}$ correlated producers who all have the same predictable behavior $b^{\mu}_{\rm prod}$.

---

[33] The dynamics in the continuum time limit reads

$$\frac{dy_i}{dt} = -\overline{\xi_i\Omega} - \sum_{j=1}^{N}\overline{\xi_i\xi_j}\tanh(y_j) + \eta\overline{\xi_i^2}\tanh y_i + \zeta_i \tag{90}$$

where ζ_i is the stochastic force. Take the average over the ensemble of realizations $y_i(\tau)$ with the same initial conditions and define $y_i(\tau) = \operatorname{arc\,tanh} m_i + \delta y_i(\tau)$, where m_i is the solution of $\min \mathcal{H}_\eta$. Then, to linear order in δy_i, for fickle agents $|m_i| < 1$, we find

$$\frac{d\langle\delta y_i\rangle}{d\tau} = -\sum_{j\ \rm fickle} T_{i,j}\langle\delta y_j\rangle + O(\delta y^2). \tag{91}$$

This dynamics become unstable when $T_{i,j}$ acquires a negative eigenvalue, i.e. when $\eta > \eta_c(\alpha)$. Note that, to this order, the dynamics is independent of Γ.

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

$$a_{\text{noise}}(t) = \text{randomsign}. \tag{95}$$

Each noise trader as a random number generator which is independent of each other agent.

It has been shown [Challet et al., 2000c, Marsili et al., 2000] 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

$$A(t) = A_{\text{spec}}(t) + A_{\text{prod}}(t) + A_{\text{noise}}(t) \tag{96}$$

where

$$A_{\text{spec}}(t) = \sum_{j=1}^{N} a_{s_j(t),j}^{\mu(t)} \tag{97}$$

and

$$A_{\text{prod}}(t) = \sum_{j=1}^{\rho N} b_j^{\mu(t)} \equiv A_{\text{prod}}^{\mu(t)} \tag{98}$$

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. 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

$$A_{\text{spec}}^\mu = \sum_{i=1}^{N} \left[a_{+1,i}^\mu \frac{1+s_i}{2} + a_{-1,i}^\mu \frac{1-s_i}{2} \right] \tag{99}$$

and A_{prod}^μ is given in Eq. (98). 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

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

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

$$G_{\text{spec}} = -\sigma^2 + \eta N - G_{\text{prod}}. \tag{100}$$

A.1 Replica calculation

The zero temperature behavior of the Hamiltonian H can be studied with spin glass techniques [Mézard et al., 1987]. We introduce n replicas of the system, each with dynamical variables $s_{i,c}$, labeled by replica indices $c,d = 1,\ldots,n$. Then we write replicated partition function:

$$\langle Z^n(\beta)\rangle = \mathrm{Tr}_s \prod_{\mu,c} \left\langle e^{-\frac{\beta}{P}(A_c^\mu)^2}\right\rangle_{a,b} \tag{101}$$

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 [Mézard et al., 1987], we introduce a Gaussian variable z_c^μ so that we can linearize the exponent in Eq. (101). 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

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

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

$$\begin{aligned} F(\hat{Q},\hat{r}) &= \frac{\alpha}{2n\beta}\mathrm{Tr}\,\log\hat{T} + \frac{\alpha\beta}{2n}\sum_{c\leq d} r_{c,d}Q_{c,d} \\ &- \frac{1}{n\beta}\log\left[\mathrm{Tr}_s e^{\frac{\alpha\beta^2}{2}\sum_{c\leq d} r_{c,d}s_c s_d}\right]. \end{aligned} \tag{102}$$

The matrix $\hat{T}$ is given by

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

For *correlated* producers we would have obtained the same result but with $\rho \to \rho + \rho^2 N\epsilon^2$, 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 $\rho \to \infty$ also corresponds to a small share of producers $\rho \ll 1$ with a small bias $\epsilon \neq 0$. Note that a bias $\epsilon \sim \sqrt{N}$ 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

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

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

$$F^{(RS)}(q,r) = \frac{\alpha}{2\beta}\log\left[1 + \frac{2(1-c)\beta(Q-q)}{\alpha}\right]$$

$$
\begin{aligned}
&+ \quad \frac{\alpha[c+\rho+(1-c)q]}{\alpha+2(1-c)\beta(Q-q)} + \frac{\alpha\beta}{2}(RQ-rq) \\
&- \quad \frac{1}{\beta}\left\langle \log \int_{-1}^{1} ds e^{-\beta V_z(s)} \right\rangle
\end{aligned} \tag{103}
$$

where we found it convenient to define the "potential"

$$
V_z(s) = -\frac{\alpha\beta(R-r)}{2} s^2 - \sqrt{\alpha r}\, z\, s \tag{104}
$$

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:

$$
\frac{\partial F^{(RS)}}{\partial q} = 0 \quad \Rightarrow \quad r = \frac{4(1-c)[c+\rho+(1-c)q]}{[\alpha+2(1-c)\beta(Q-q)]^2} \tag{105}
$$

$$
\frac{\partial F^{(RS)}}{\partial Q} = 0 \quad \Rightarrow \quad \beta(R-r) = -\frac{2(1-c)}{\alpha+2(1-c)\beta(Q-q)} \tag{106}
$$

$$
\frac{\partial F^{(RS)}}{\partial R} = 0 \quad \Rightarrow \quad Q = \langle\langle s^2 \rangle\rangle \tag{107}
$$

$$
\frac{\partial F^{(RS)}}{\partial r} = 0 \quad \Rightarrow \quad \beta(Q-q) = \frac{\langle\langle sz \rangle\rangle}{\sqrt{\alpha r}} \tag{108}
$$

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

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

$$
\chi = \frac{2(1-c)\beta(Q-q)}{\alpha}, \quad \text{and} \quad \zeta = -\sqrt{\frac{\alpha}{r}}\beta(R-r) \tag{109}
$$

and to require that they stay finite in the limit $\beta \to \infty$. 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 $z \le -\zeta$ and at $s = +1$ for $z \ge \zeta$. For $-\zeta < z < \zeta$ the minimum is at $s = z/\zeta$. With this we find

$$
\langle\langle sz \rangle\rangle = \frac{1}{\zeta}\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) \tag{110}
$$

and

$$
\langle\langle s^2 \rangle\rangle = Q = 1 - \sqrt{\frac{2}{\pi}}\frac{e^{-\zeta^2/2}}{\zeta} - \left(1 - \frac{1}{\zeta^2}\right)\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) \tag{111}
$$

With some more algebra, one easily finds:

$$
\chi = \left[\alpha/\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) - 1\right]^{-1} \tag{112}
$$

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

$$
\sqrt{\frac{2}{\pi}}\frac{e^{-\zeta^2/2}}{\zeta} + \left(1 - \frac{1}{\zeta^2}\right)\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) + \frac{\alpha}{\zeta^2} = \frac{1+\rho}{1-c} \tag{113}
$$

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 $c \to 1$, i.e. when speculators become producers, or if $\rho \to \infty$ (many producers).

Eq. (112) means that χ diverges when $\alpha \to \alpha_c(\rho, c)^+$, which then implies that at the critical point

$$\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) = \alpha = \alpha_c. \tag{114}$$

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

$$\sqrt{\frac{2}{\pi}}\frac{e^{-\zeta^2/2}}{\zeta} + \mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) = \frac{1+\rho}{1-c}. \tag{115}$$

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

$$F^{(RS)} = \frac{c + (1-c)Q + \rho}{(1+\chi)^2} \tag{116}$$

where Q and χ take their saddle point values Eqs. (111) and (112).

The gain of producers, from Eq. (103), is

$$\frac{G_{\mathrm{prod}}}{N} = -\frac{\rho}{1+\chi} \tag{117}$$

and that of speculators is obtained from Eq. (100).

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

The phase diagram is shown in figure 18. More details on the behavior of the solution are discussed in Ref. [Challet et al., 2000d].

References

[Arthur, 1994] Arthur, B. W. (1994). Inductive reasoning and bounded rationality: the el farol problem. *Am. Econ. Rev.*, 84:406–411.

[Arthur et al., 1997] Arthur, W. B., Holland, J. H., LeBaron, B., Palmer, R., and Tayler, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. In Arthur, W. B., Durlauf, S., and Lane, D., editors, *The Economy as an Evolving Complex System, II*, volume XXVII of *SFI Studies in the Sciences of Complexit*. Addison-Wesley.

[Bottazzi et al., 2003] Bottazzi, G., G.Devetag, and G.Dosi (2003). *Heterogeneous Agents, Interactions and Economic Performance*, chapter Adaptive Learning and Emergent Coordination in Minority Games. Number 521 in Lecture Notes in Economics and Mathematical Systems. Springer Verlag.

[Bouchaud and Potters, 2000] Bouchaud, J.-P. and Potters, M. (2000). *Theory of Financial Risks*. Cambridge University Press, Cambridge.

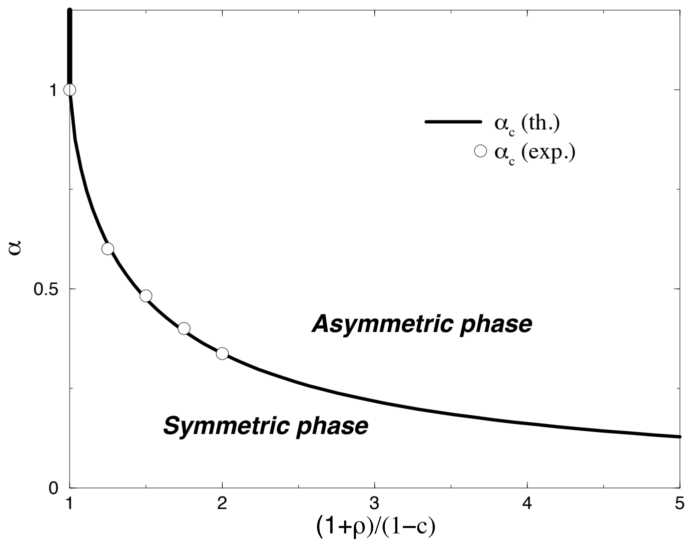

Figure 18: Phase diagram $\alpha_c[(1+\rho)/(1-c)]$

[Caldarelli et al., 1997] Caldarelli, G., Marsili, M., and Zhang, Y.-C. (1997). A prototype model of stock exchange. *Europhysics Letters*, 50:479–484.

[Caridi and Ceva, 2003] Caridi, I. and Ceva, H. (2003). Minority game: a mean-field-like approach. *Physica A*, 317:247–258.

[Cavagna, 1999] Cavagna, A. (1999). Irrelevance of memory in the minority game. *Phys. Rev. E*, 59:R3783–R3786. cond-mat/9812215.

[Cavagna et al., 1999] Cavagna, A. et al. (1999). A thermal model for adaptive competition in a market. *Phys. Rev. Lett.*, 83:4429–4432. cond-mat/9903415.

[Cavagna et al., 2000] Cavagna, A., Garrahan, J. P., Giardina, I., and Sherrington, D. (2000). Reply to comment on "thermal model for adaptive competition in a market". *Phys. Rev. Lett.*, 85:5009. cond-mat/9905134.

[Challet,] Challet, D. The minority game's web page `www.unifr.ch/econophysics/minority`.

[Challet, 2003] Challet, D. (2003). *Competition between adaptive agents: from learning to collective efficiency and back*. Springler. cond-mat/0210319.

[Challet et al., 2000a] Challet, D., Chessa, A., Marsili, M., and Zhang, Y.-C. (2000a). From minority games to real markets. *Quant. Fin.*, 1:168. cond-mat/0011042.

[Challet and Marsili, 1999] Challet, D. and Marsili, M. (1999). Symmetry breaking and phase transition in the minority game. *Phys Rev. E*, 60:R6271. cond-mat/9904392.

[Challet and Marsili, 2000] Challet, D. and Marsili, M. (2000). Relevance of memory in minority games. *Phys. Rev. E*, 62:1862. cond-mat/0004196.

[Challet and Marsili, 2003a] Challet, D. and Marsili, M. (2003a). Criticality and finite size effects in a simple realistic model of stock market.

[Challet and Marsili, 2003b] Challet, D. and Marsili, M. (2003b). Shedding light on el farol. *Physica A*.

[Challet et al., 2000b] Challet, D., Marsili, M., and Zecchina, R. (2000b). Comment on "thermal model for adaptive competition in a market". *Phys. Rev. Lett.*, 85:5008. cond-mat/9904392.

[Challet et al., 2000c] Challet, D., Marsili, M., and Zecchina, R. (2000c). Statistical mechanics of heterogeneous agents: minority games. *Phys. Rev. Lett.*, 84:1824–1827. cond-mat/9904392.

[Challet et al., 2000d] Challet, D., Marsili, M., and Zhang, Y.-C. (2000d). Modeling market mechanisms with minority game. *Physica A*, 276:284. cond-mat/9909265.

[Challet et al., 2001] Challet, D., Marsili, M., and Zhang, Y.-C. (2001). Stylized facts of financial markets in minority games. *Physica A*, 294:514. cond-mat/0101326.

[Challet and Zhang, 1997] Challet, D. and Zhang, Y.-C. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407. adap-org/9708006.

[Challet and Zhang, 1998] Challet, D. and Zhang, Y.-C. (1998). On the minority game: analytical and numerical studies. *Physica A*, 256:514. cond-mat/9805084.

[Cont and Bouchaud, 2000] Cont, R. and Bouchaud, J.-P. (2000). Herd behaviour and aggregate fluctuation in financial markets. *Macroecon. Dyn.*, 4:170.

[Coolen and Heimel, 2001] Coolen, A. A. C. and Heimel, J. A. F. (2001). Dynamical solution of the on-line minority game. *J. Phys. A: Math. Gen.*, 34:10783–10804.

[de Almeida and Thouless, 1978] de Almeida, J. and Thouless, D. (1978). Stability of the sherrington-kirkpatrick solution of a spin glass model. *JPHA*, 11:983.

[D'hulst and Rodgers, 1999] D'hulst, R. and Rodgers, G. (1999). The hamming distance in the minority game. *Physica A*, 270:514. cond-mat/9902001.

[Eissfeller and Opper, 1992] Eissfeller, H. and Opper, M. (1992). New method for studying the dynamics of disordered spin systems without finite-size effects. *Phys. Rev. Lett.*, 68.

[Farmer, 1999] Farmer, J. D. (1999). Physicists attempt to scale the ivory towers of finance. *Comp. Fin.*, pages 26–39.

[Farmer and Joshi, 99] Farmer, J. D. and Joshi, S. (99).

[Fudenberg and Tirole, 1991] Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press.

[Galambos, 1987] Galambos, J. (1987). *Asymptotic theory of extreme order statistics*. Krieger, Malabar, Florida, 2 edition.

[Garrahan et al., 2000] Garrahan, J. P., Moro, E., and Sherrington, D. (2000). Continuous time dynamics of the thermal minority game. *Phys. Rev. E*, 62:R9. cond-mat/0004277.

[Giardina and Bouchaud, 2002] Giardina, I. and Bouchaud, J.-P. (2002). Crashes and intermittency in agent based market models.

[Hart et al., 2001] Hart, M., Jefferies, P., Hui, P., and Johnson, N. (2001). Crowd-anticrowd theory of multi-agent market games. *Eur. Phys. J. B*, 20:547–550. cond-mat/0008385.

[Hart et al., 2002] Hart, M. L., Jefferies, P., and Johnson, N. F. (2002). Dynamics of the time horizon minority game. *Physica A*, 311:275. cond-mat/0103259.

[Heimel and Coolen, 2001] Heimel, J. A. F. and Coolen, A. A. C. (2001). Generating functional analysis of the dynamics of the batch minority game with random external information. *Phys. Rev. E*, 63:056121. cond-mat/0012045.

[Heimel et al., 2001] Heimel, J. A. F., Coolen, A. A. C., and Sherrington, D. (2001). Dynamics of the batch minority game with inhomogeneous decision noise. *Phys. Rev. E*, 65:016126. cond-mat/0106635.

[Heimel and Martino, 2001] Heimel, J. A. F. and Martino, A. D. (2001). Broken ergodicity and memory in the minority game. *J. Phys. A: Math. Gen.*, 34:L539–L545. cond-mat/0108066.

[Hod and Nakar, 2002] Hod, S. and Nakar, E. (2002). Self-segregation versus clustering in the evolutionary minority game. *Phys. Rev. Lett.*, 88:238702.

[Hui et al., 1999] Hui, N. J. P., Zheng, D., and Hart, M. (1999). Enhanced winnings in a mixed-ability population playing a minority game. *J. Phys. A: Math. Gen.*, 32:L427.

[Johnson et al., 1999a] Johnson, N., Hart, M., and Hui, P. (1999a). Crowd effects and volatility in a competitive market. *Physica A*, 269:1. cond-mat/9811227.

[Johnson et al., 1999b] Johnson, N., Hui, P., Zheng, D., and Tai, C. (1999b). Minority game with arbitrary cutoffs. *Physica A*, 269:493.

[Johnson et al., 1998] Johnson, N., Jarvis, S., Jonson, R., Cheung, P., Kwong, Y., and Hui, P. (1998). Volatility and agent adaptability in a self-organizing market. *Physica A*, 256:230. cond-mat/9802177.

[Kinzel et al., 2000] Kinzel, W., Metzler, R., and Kanter, I. (2000). Dynamics of interacting neural networks. *J. Phys. A: Math. Gen.*, 33:L141–147. cond-mat/9906058.

[Lamper et al., 2002] Lamper, D., Howison, S., and Johnson, N. (2002). Predictability of large future changes in a competitive evolving population. *Phys. Rev. Lett.*, 88:017902–017905. cond-mat/0105258.

[Levy et al., 2000] Levy, H., Levy, M., and Solomon, S. (2000). *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena.* Academic Press, London.

[Li et al., 2000] Li, Y., A., A. V., and Savit, R. (2000). The minority game with variable payoffs. *Physica A*, 284:461–477. cond-mat/0002004.

[Luce, 1959] Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis.* Wiley.

[Lux and Marchesi, 1999] Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:498–500.

[Mansilla, 2000] Mansilla, R. (2000). From naive to sophisticated behavior in multiagents based financial markets models. *Physica A*, 284:478–488. cond-mat/0002331.

[Mantegna and Stanley, 2000] Mantegna, R. and Stanley, H. G. (2000). *Introduction to Econophysics.* Cambridge University Press.

[Manuca et al., 2000] Manuca, R., Li, Y., Riolo, R., and Savit, R. (2000). The structure of adaptive competition in minority games. *Physica A*, 282:559–608.

[Marsili, 2001] Marsili, M. (2001). Market mechanism and expectations in minority and majority games. *Physica A*, 299:93–103.

[Marsili and Challet, 2001a] Marsili, M. and Challet, D. (2001a). Continuum time limit and stationary states of the minority game. *Phys. Rev. E*, 64:056138. cond-mat/0102257.

[Marsili and Challet, 2001b] Marsili, M. and Challet, D. (2001b). Trading behavior and excess volatility in toy markets. *Adv. Complex Systems*, 3(I):3–17. cond-mat/0011042.

[Marsili et al., 2000] Marsili, M., Challet, D., and Zecchina, R. (2000). Exact solution of a modified el farol's bar problem: Efficiency and the role of market impact. *Physica A*, 280:522. cond-mat/9908480.

[Marsili et al., 2001] Marsili, M., Mulet, R., Ricci-Tersenghi, F., and Zecchina, R. (2001). Learning to coordinate in a complex and nonstationary world. *Phys. Rev. Lett.*, 87:208701. cond-mat/0105345.

[Marsili and Piai, 2002] Marsili, M. and Piai, M. (2002). Colored minority games. *Physica A*, 310:234–244. cond-mat/0202479.

[Martino and Marsili, 2001] Martino, A. D. and Marsili, M. (2001). Replica symmetry breaking in the minority game. *J. Phys. A: Math. Gen.*, 34:2525–2537. cond-mat/0007397.

[Martino et al., 2003] Martino, A. D., Marsili, M., and Mulet, R. (2003). Adaptive drivers in a model of urban traffic. cond-mat/0308543.

[McFadden, 1981] McFadden, D. (1981). Econometric models of probabilistic choice. In Manski, C. F. and McFadden, D., editors, *Structural analysis of discrete data with econometric application*, pages 171–260. MIT Press.

[Metzler, 2002] Metzler, R. (2002). Antipersistent binary time series. *J. Phys. A: Math. Gen.*, 35:721–730.

[Mézard et al., 1987] Mézard, M., Parisi, G., and Virasoro, M. A. (1987). *Spin glass theory and beyond*. World Scientific.

[P. Jefferies, 2001] P. Jefferies, M.L. Hart, N. J. (2001). Deterministic dynamics in the minority game. *Phys. Rev. E*, 65:016105. cond-mat/0102384.

[Paczuski et al., 2000] Paczuski, M., Bassler, K. E., and Corral, A. (2000). Self-organized networks of competing boolean agents. *Phys. Rev. Lett.*, 84(14):3185–3188. cond-mat/9905082.

[Reents et al., 2001] Reents, G., Metzler, R., and Kinzel, W. (2001). A stochastic strategy for the minority game. *Physica A*, 299:253–261. cond-mat/0007351.

[Rustichini, 1999] Rustichini, A. (1999). Optimal properties of stimulus response models. *Games and Economic Behavior*, 29.

[Savit et al., 1999] Savit, R., Manuca, R., and Riolo, R. (1999). Adaptive competition, market efficiency, and phase transitions. *Phys. Rev. Lett.*, 82:2203–2206.

[Sengupta and Mitra, 1999] Sengupta, A. M. and Mitra, P. P. (1999). Distributions of singular values for some random matrices. *Phys. Rev. E*, 60:3389. cond-mat/9709283.

[Sherrington and Galla, 2003] Sherrington, D. and Galla, T. (2003). The minority game: effects of strategy correlations and timing of adaptation. *Physica A*, 324:25–29.

[Sherrington et al., 2002] Sherrington, D., Moro, E., and Garrahan, J. P. (2002). Statistical physics of induced correlation in a simple market. *Physica A*, 311:527–535. cond-mat/0010455.

[Wakeling and Bak, 2001] Wakeling, J. and Bak, P. (2001). Intelligent systems in the context of surrounding environment. *Phys. Rev. E*, 64:051920. nlin.AO/0201046.

[Yeomans, 1992] Yeomans, J. (1992). *Statistical Mechanics of Phase Transitions*. Oxford University Press.

[Zhang, 1998] Zhang, Y.-C. (1998). Modeling market mechanism with evolutionary games. *Europhys. News*, 29:51.

[Zhang, 1999] Zhang, Y.-C. (1999). Towards a theory of marginally efficient markets. *Physica A*, 269:30.