SISSA ISAS

SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI
INTERNATIONAL SCHOOL FOR ADVANCED STUDIES

Replica-symmetry breaking and long-term memory in large games with heterogeneous players

Thesis submitted for the degree of
"Doctor Philosophiæ"

CANDIDATE

Andrea De Martino

SUPERVISOR

Dr Matteo Marsili

Trieste, September 2001

Summary

Mean-field spin-glasses at low temperatures display complex equilibrium and off-equilibrium properties, whose theoretical description relies heavily on the concepts of replica-symmetry breaking, explaining the complicated phase-space structure, and weak ergodicity breaking, accounting for fact that glassy systems never equilibrate in any finite region of phase space. The key assumption that makes the analysis of the persistent out-of-equilibrium (or aging) regime possible is the absence of long-term memory: the system for long times forgets the initial transients. Consequently, the dynamical equations governing the behaviour of correlation and response functions can be solved asymptotically by gauging away all uncontrollable short-time details. Aging manifests with the existence of asymptotic solutions violating time-translation invariance.

In this thesis we study the minority game, a model defined by a set of microscopic stochastic equations schematizing the behaviour of speculators in a financial market. Our aim is to characterize the collective behaviour by analyzing the stationary states. The system lacks of an explicit global Hamiltonian. However, the dynamical steady states are found to correspond to the local minima of a disordered mean-field spin-system. Therefore, the toolkit of spin-glass theory (the replica method and the dynamical functional-integral technique) applies.

It is shown that the system undergoes a static replica-symmetry breaking transition completely analogous to that occurring in spin-glasses. However, the dynamical scenario in the non-ergodic regime turns out to be quite different from the usual. In particular, we show that ergodicity breaking in this case is related to the onset of long-term memory. For any initial condition, the system reaches a stationary state and, consequently, no aging takes place (time-translation invariance is always satisfied) and ergodicity breaking is 'strong'. We are unfortunately unable to characterize non-ergodic stationary states, due to the presence of long-term memory. However, we derive a condition for the breakdown of the weak memory assumption, by which we solve the ergodic regime. Such a condition is found to be identical to the static AT-line, which marks the onset of replica-symmetry breaking. Static and dynamic results agree perfectly, and are supported by extensive computer simulations. We discuss this non-standard picture and propose a physical explanation.

Finally, the statics of two economics-inspired models obtained as variations on minority game's theme is studied and discuss.

Acknowledgments. This work has greatly benefited from discussions I had with several people: F. Cecconi, D. Challet, A. C. C. Coolen, A. Engel, A. Flammini, S. Franz, L. Giada, A. Giansanti, J. A. F. Heimel, G. Lattanzi, M. Leone, D. Marenduzzo, F. Ricci-Tersenghi, A. Trombettoni, and R. Zecchina. I also wish to acknowledge The Abdus Salam ICTP, King's College London and the Max-Planck-Institut für Physik Komplexer Systeme for hospitality and support at different stages of this work. Finally, I want to thank my boss Matteo.

Contents

CHAPTER 1

Introduction

Mean-field models of spin-glasses possess highly non-trivial static and dynamic properties [1–4]. The low-temperature phase is characterized by the existence of a huge number of free-energy minima separated by barriers, some of which diverge with the system size. This causes ergodicity breaking in the thermodynamic limit: as the system cannot override infinitely high barriers, its time evolution is restricted to a single valley in the free-energy hyper-surface, corresponding to an isolated ergodic component. Such a valley is however still rugged, with (infinitely many, in the thermodynamic limit) barriers of finite heights separating metastable states. In such a context, the relaxational dynamics starting from a random (or high-temperature) configuration becomes exceedingly slow, because transitions between metastable states stretch equilibration times. The most striking signature of this persistent out-of-equilibrium regime is aging, namely the fact that the system relaxes more slowly the older it is, i.e. the longer the waiting time t_w elapsed between the quench from high temperatures, when the dynamics starts ($t = 0$), and the beginning of measurements. Aging has a simple physical explanation. During their time evolution, systems cross low free-energy barriers first, since these can be easily surmounted. For this reason a young system (small t_w) relaxes faster. Older systems, instead, having already crossed the low barriers, are confronting with the high ones, which are much harder to overcome. Moreover, the distribution of the heights of the barriers within a valley has no upper bound, i.e. there are finite barriers of all heights. As a consequence, equilibrium is never achieved: there is no time t_{eq} after which the configurations are visited with a Boltzmann-Gibbs probability distribution. More precisely, the system does not equilibrate within any finite region of the phase space, because it does not remain trapped in any metastable state indefinitely. All this can be seen in the behaviour of two-time quantities (macroscopic observables that depend on the configuration of the system at two different times), like correlation and response functions. In contrast with the ergodic regime, where they tend to some stationary value, at low temperatures they remain asymptotically time- and history-dependent.

The analytic study of such models turned out to be a formidable problem. In fact, standard statistical-mechanics techniques to evaluate the partition function are totally inadequate for these systems. The problem is that the ground states of the disordered Hamiltonian are not related to each other by any (trivial) symmetry transformation, e.g. spin-flip, so that it is not possible to select a single equilibrium state by applying a suitably chosen external field. Moreover, as for physical reasons it is to be expected that the relevant thermodynamic functionals do not depend on the specific realization of disorder in the infinite-volume limit, one usually has to calculate disorder-averaged (or "typical") quantities, which is in principle a very difficult task. However, in the past years a number of powerful techniques have been developed to overcome these problems, and presently both the statics and the dynamics are analytically tractable.

Equilibrium theory relies on the so-called replica trick, which allows to manage disorder. It consists essentially in evaluating the partition function of a new system of n identical non-interacting copies, or replicas, of the original system. This new system turns out to encompass all the physics of the original model. By construction, the new Hamiltonian is invariant under permutation of the n replicas. High-temperature properties are described by solutions preserving this symmetry. At low temperatures, where the phase space is fragmented into many separated ergodic components, the correct physics is instead given by solutions with broken replica-permutation symmetry. As usual in phase transitions with symmetry breaking, the original symmetry group, i.e. the group of permutations of n elements, reduces to one of its subgroups. The celebrated Parisi scheme of replica-symmetry breaking expresses this subgroup by providing the physically-significant pattern of symmetry breaking.

The dynamical counterpart of the replica method is the dynamical functional-integral technique [5,6] (also called generating functional analysis or dynamical mean-field theory) à la De Dominicis [7]. It consists in introducing a functional that generates the relevant macroscopic observables, i.e. the disorder-averaged correlation and response functions, as suitable derivatives. Starting from Markovian equations of motion for the microscopic variables, one finds that, in the infinite-volume limit, these quantities evolve according to non-Markovian stochastic equations. In the high-temperature regime, the physical solutions are time-translation invariant (i.e. homogeneous in time), and related by the fluctuation-dissipation theorem. At low temperatures, instead, solutions breaking time-translation invariance appear. This signals the onset of the aging regime. What really makes the dynamics in the non-ergodic phase tractable is the absence of long-term memory. By this one means essentially that any perturbation lasting for a finite time will be forgotten in the long run, so that the dynamics for long times bears virtually no memory of the initial transients. This property, which has been termed "weak long-term memory" in the literature [8], is trivially verified at high temperatures, and remains valid also at low-temperatures, where time-translation invariance fails.

The physical relevance of these studies lies in the guess that the complexity of mean-field models survives in finite-dimensional models , and is adequate for describing the complex behaviour of real systems (see e.g. [9]). Indeed, what has clearly emerged from analytic theories and computer experiments is that mean-field models of spin-glasses do capture some of the non-equilibrium features observed in real spin-glasses (especially models with a continuous static transition) and in structural glasses (especially models with a discontinuous static transition) [4]. Nevertheless, the infinite-connectivity assumption on which mean-field theory is based is hard to justify for real physical systems. In other cases, however, and particularly in applications outside physics, it turns out to be more appropriate. Famous examples are found in combinatorial optimization [10] and neural networks [11]. The common solution strategy of all these problems consists in mapping them onto a mean-field spin-glass, for which a full abstract theoretical framework and a deal of technical tools are readily available. The mean-field theory of spin-glasses has therefore come to play a significant unifying role across several branches of science, and its picture of *complexity* has been extended to many complex systems outside physics.

One of the most recently-discovered sources of inspiration and challenges for the spin-glass paradigm is found in economics. Systems such as financial markets definitely qualify as complex [12]. From a theoretical viewpoint, they seem to be suited for a statistical description (provided one accepts to deal with large systems only). If the complicated behaviour of the individual traders is modeled by sufficiently simple stochastic microscopic laws, it is possible to derive equations for some relevant macroscopic observables, like the price of an asset or the magnitude of price fluctuations. Furthermore, it is quite reasonable to assume that traders in a financial market make their trading decisions by looking at some macroscopic observable, say the price. Under this hypothesis, microscopic interactions among economic agents are mediated by a macroscopic quantity which they all contribute to form, and the situation is equivalent to postulating a mean-field type of microscopic interaction. Contact with spin-glasses is made if one introduces two further ingredients, namely heterogeneity of agents and limited availability of resources. The former refers to the fact that, in principle, each agent reacts differently to the receipt of exogenous (e.g. political news or the moon-phases) or endogenous (e.g. price changes) information, and translates into the presence of quenched or annealed disorder. The latter causes frustration by making the traders' desires not all simultaneously satisfiable. The toolkit of spin-glass theory becomes at this point indispensable, and the hope is to find in the spin-glass paradigm an explanation for the complex features observed in real markets.

The minority game [13] is a model which aims to capture the collective phenomena occurring in systems of heterogeneous interacting agents within the simplest non-trivial mathematical framework. It is a variation on the so-called El-Farol problem [14]. One considers a system of N agents, with N large. At each point in time, all agents receive an information pattern and, basing on this, each of them has to make a binary decision (buy/sell). Those who end up in the minority side have a profit, while the others face a loss. This minority-wins mechanism is basically an abstraction of the law of supply-and-demand: when the majority of traders is buying it is convenient to sell since prices are likely to be high, and vice versa. Making it impossible for all traders to be successful at the same time, such a rule provides a high degree of frustration. At the microscopic level, agents are assumed to react to an endogenous macroscopic quantity that changes over time (specifically, it is the difference between the number of buyers and the number of sellers, which may be likened to the "excess demand" in economics), which acts essentially as a time-dependent molecular field. This microscopic dynamics describes a "learning" behaviour, with which agents try to optimize their decision-making process in time. Quenched disorder derives from the different ways in which agents react to the receipt of external information patterns. In the simplest setup, the information is assumed to be randomly drawn at each time step from a set of P possible patterns.

The central problem is to analyze the stationary state (if any) of the game, which is repeated at each time step, in order to grasp the conditions upon which an (in a precise sense) efficient state is achievable. The phenomenology one observes in computer simulations is remarkably rich. Upon varying the relative number $\alpha = P/N$ of information patterns, one finds, among other things, a transition from an ergodic regime (high α), where a unique stationary state is present, to a highly non-ergodic regime (low α), where the stationary state depends strongly on the initial conditions. In the non-ergodic regime, in particular, both very efficient and very inefficient states can occur.

A detailed understanding of this complex behaviour has required a considerable effort. The major technical difficulty lies in the fact that the model is not based on an explicit Hamiltonian. In this respect, it is not evident that one can study the stationary states by equilibrium techniques. However, following a few years of intense research, a static theory has been developed by mapping the steady states of the dynamics onto the minima of a mean-field spin-glass Hamiltonian H_0 [15–21], and the dynamics has been solved exactly via the dynamical functional-integral formalism [22–24], for which to apply the existence of a global Hamiltonian is not a prerequisite. The static theory in particular has been the subject of much debate. The problem is that the minima of H_0 are always given by replica-symmetric solutions, meaning that in principle it describes adequately only ergodic stationary states (and this conclusion is confirmed by numerical simulations). When ergodicity is broken, the behaviour of the system is of a purely dynamical nature, and its connection to the static replica-symmetric description is quite subtle [25].

In this thesis we study a generalized version of the minority game. The crucial difference lies in the fact that we add an Onsager reaction field to the molecular field "felt" by agents. This new term describes the contribution of agent i himself to the molecular field. Our choice is motivated by numerical results [17, 18] suggesting a substantial phenomenological enrichment. We modulate the strength of the reaction field, and thus its impact on the dynamics, with a real number η, so that for $\eta = 0$, the original setup is retrieved, and by tuning η one can describe different types of economic agents.

In this modified setup, an explicit connection between broken ergodicity and static replica-symmetry breaking occurs. This allows in principle to formulate a closer parallel with the spin-glass paradigm. However, replica-symmetry breaking turns out not to be accompanied by aging: correlation and response functions remain time-translation invariant even in the non-ergodic regime. Rather, we find that non-ergodicity in this system is tightly related to the onset of long-term memory. Strikingly, the boundary between the replica-symmetric phase and the phase with broken replica-symmetry (the so-called AT-line) coincides with the condition for the onset of memory. This picture, which is supported by computer experiments, is quite unusual for disordered systems. We will discuss a possible physical reason for the origin of this behaviour and indicate a class of statistical models, with a necessarily dynamical definition, that are likely to behave in the same manner. We stress the fact that the study of systems whose memory is not 'weak' is, to our knowledge, a completely open issue.

The work is organized as follows. In Chapter 2 we shall define our model, set the notation, describe briefly the techniques employed, and summarize the properties of the original minority game. Chapter 3 deals with the equilibrium properties of the generalized model [26]. We move to the dynamics in Chapter 4, where we also compare our findings with the usual scenario [27]. In Chapter 5 we present two economics-inspired variations of the model [28], which aim to discuss, respectively, the behaviour of a particular class of economic agents and the origin of market organization. Finally, in Chapter 6, we formulate our conclusions.

CHAPTER 2

The minority game

This chapter is dedicated to a preliminary discussion of the generalized minority game. First, we shall define the model and its mathematical structure. Then, we will briefly describe the analytical techniques that will be employed in its study. Finally, the properties observed in numerical simulations and the theory of the original model will be reviewed.

2.1. Definition of the model

We consider N agents (or players) labeled by roman indices like i and j. N is supposed to be very large. Eventually, the limit $N \to \infty$ will be examined. The external information patterns, labeled by μ, can take on extensively many different values, so that $\mu \in \{1, \ldots, \alpha N\}$. At each iteration round $n = 1, 2, \ldots$ an information pattern $\mu(n)$ is chosen at random from $\{1, \ldots, \alpha N\}$ with uniform probabilities and reported to players. All agents receive the same information $\mu(n)$. Each agent has at his disposal S different strategies (labeled by g) to convert the acquired information into a trading decision. Strategies are denoted by αN-dimensional vectors:

$$\boldsymbol{a}_{ig} = \{a_{ig}^{\mu}\}_{1 \le \mu \le \alpha N} \in \{-1, 1\}^{\alpha N} \qquad (1 \le i \le N \ ; \ 1 \le g \le S) \tag{2.1}$$

$a_{ig}^{\mu} \in \{-1, 1\}$ should be interpreted as the trading action (e.g. 'buy' and 'sell' for -1 and 1, respectively) prescribed to agent i by his g-th strategy given receipt of information μ. By assumption, each component a_{ig}^{μ} is selected randomly and independently from $\{-1, 1\}$ with uniform probabilities before the start of the game for all i, g and μ, and is kept fixed throughout the game. This is the source of quenched disorder.

Each strategy of every agent is given an initial valuation $p_{ig}(0)$, which is updated at the end of every round. The valuation $p_{ig}(n)$ measures, roughly speaking, the perceived success of strategy g up to round n. At the start of round n, given $\mu(n)$, every agent selects the strategy with the highest valuation, i.e. $\boldsymbol{a}_{i\widetilde{g}_i(n)}$ with

$$\widetilde{g}_i(n) = \arg\max\, p_{ig}(n) \tag{2.2}$$

and subsequently makes a bid $b_i(n) \in \{-1, 1\}$ according to the trading decision set by the selected strategy, $b_i(n) = a_{i\widetilde{g}_i(n)}^{\mu(n)}$. The total bid at round n is given by

$$A(n) \equiv A^{\mu(n)} = \sum_{1 \le i \le N} b_i(n) \tag{2.3}$$

Notice that $A(n) = \mathcal{O}(\sqrt{N})$. Finally, for all i and g all valuations are updated according to the reinforcement learning dynamics

$$p_{ig}(n+1) = p_{ig}(n) - \frac{1}{N}\, a_{ig}^{\mu(n)} \left[A(n) - \eta \left(a_{i\widetilde{g}_i(n)}^{\mu(n)} - a_{ig}^{\mu(n)}\right)\right] \tag{2.4}$$

and agents move into the next round.

Let us discuss the meaning of (2.4) starting from the case $\eta = 0$, to which the original definition of Challet and Zhang [13] corresponds. One had simply $p_{ig}(n+1) = p_{ig}(n) - a_{ig}^{\mu(n)} A(n)/N$, so that at each stage of the game all those strategies that would have produced a minority decision *given the actually occurred* $A(n)$ (i.e. those $\boldsymbol{a}_{ig}$ for which $a_{ig}^{\mu(n)} A(n) < 0$) are rewarded. In this setup agents completely neglect the fact that if they had used strategy $\boldsymbol{a}_{ig}$ instead of $\boldsymbol{a}_{i\widetilde{g}_i(n)}$ their contribution to $A(n)$, and thus $A(n)$ itself, could have changed. The term proportional to η in square brackets, which we call *market-impact correction*, is non-zero only for $g \neq \widetilde{g}_i(n)$ and adjusts the updating rule by estimating the eventual change in the total bid induced by choosing g instead of $\widetilde{g}_i(n)$. It acts essentially like an Onsager reaction field [2]. The parameter η sets the strength of the adjustment, and allows to interpolate smoothly between the original model ($\eta = 0$) and the case in which the total bid is completely adjusted ($\eta = 1$).

Of course, η can be negative or larger than 1 as well An alternative description of the dynamics might help to clarify the agents' behaviour in these cases. It stems from the following remark. The correction to the $\eta = 0$ dynamics is given by $\eta(a^{\mu(n)}_{i\widetilde{g}_i(n)}a^{\mu(n)}_{ig} - 1)/N$. Observing that $a^{\mu(n)}_{i\widetilde{g}_i(n)}a^{\mu(n)}_{ig}$ is 1 when $\widetilde{g}_i(n) = g$ and a random sign, with zero average, otherwise, one may loosely argue that "on the average" the correction is $\eta(\delta_{\widetilde{g}_i(n),g} - 1)/N$, where $\delta_{i,j}$ is the Kronecker δ-symbol. From this it is clear that when $\eta > 0$ agents over-reward the actually played strategy by a quantity proportional to η, while for $\eta < 0$ they under-reward it by a quantity proportional to $|\eta|$.

In economic terms, the model with $\eta = 1$ describes agents with full information, who not only know the total bid that actually occurred at round n (with them using strategy $\boldsymbol{a}_{i\widetilde{g}_i(n)}$), but also know exactly what the total bid would have been had they followed a different strategy. In game theoretic language, they are *sophisticated agents*. The case $\eta = 0$, instead, describes *naive agents*, whose knowledge of the situation is only partial. As they can't grasp their contribution to $A(n)$, they assume that it wouldn't have changed if they acted differently. In economic terms, one says they behave as *price-takers*, not taking their own contribution to the total bid into any account [18].

The key macroscopic observables in the model are the time average of the total bid in the stationary state of the process (supposedly reached after ℓ rounds), and its variance, called *volatility*, defined respectively as

$$\langle A\rangle_{\text{time}} = \lim_{L\to\infty}\frac{1}{L-\ell}\sum_{\ell\le n\le L} A(n) \qquad \text{and} \qquad \sigma^2 = \left\langle A^2\right\rangle_{\text{time}} - \langle A\rangle^2_{\text{time}} \tag{2.5}$$

The former measures the average mismatch between buyers and sellers (*excess demand*), while the latter describes magnitude of market fluctuations. Notice that the re-scaled quantities $\langle A\rangle_{\text{time}}/\sqrt{N}$ and σ^2/N remain bounded as $N \to \infty$. Moreover, if the limit $L \to \infty$ is taken before the thermodynamic limit $N \to \infty$, then both quantities will be self-averaging with respect to both the realization of the quenched disorder $\{\boldsymbol{a}_{ig}\}$ and the realization of the sequence of information patterns $\{\mu(n)\}_{n\le L}$. These properties makes them suitable macroscopic observables for statistical analysis.

Efficient systems are defined as having zero excess demand and "small" volatility. In our case, it turns out that indeed $\langle A\rangle_{\text{time}} = 0$ (by symmetry), so that we will concentrate on the volatility

$$\sigma^2 = \left\langle A^2\right\rangle_{\text{time}} = \lim_{L\to\infty}\frac{1}{L-\ell}\sum_{\ell\le n\le L} A(n)^2 \tag{2.6}$$

In order to characterize "small", let us notice that random trading, where agents decide their actions by coin-tossing, implies $\langle A\rangle_{\text{time}} = 0$ and $\sigma^2/N = 1$. We will use the latter as a reference value, and call a system *efficient* when $\sigma^2/N < 1$. As the dependence of the volatility on η is crucial in this work, we will always write it explicitly as σ^2_η.

In the following, without any loss of generality, we will concentrate on the case $S = 2$, where equations simplify considerably upon introducing, for each i, the quantities

$$y_i(n) = \frac{1}{2}\left[p_{i1}(n) - p_{i2}(n)\right] \qquad \boldsymbol{\omega}_i = \frac{1}{2}\left(\boldsymbol{a}_{i1} + \boldsymbol{a}_{i2}\right) \qquad \boldsymbol{\xi}_i = \frac{1}{2}\left(\boldsymbol{a}_{i1} - \boldsymbol{a}_{i2}\right) \tag{2.7}$$

We call the variable $y_i(n)$ the *preference* of agent i. It allows to re-write the strategy selected by agent i at round n and the total bid $A(n)$ respectively as

$$\boldsymbol{a}_{i\widetilde{g}_i(n)} = \boldsymbol{\omega}_i + \operatorname{sign}\left[y_i(n)\right]\boldsymbol{\xi}_i \tag{2.8}$$

$$A(n) \equiv A^{\mu(n)} = \Omega^{\mu(n)} + \sum_{1\le j\le N} \xi_i^{\mu(n)}\operatorname{sign}\left[y_i(n)\right] \tag{2.9}$$

with $\boldsymbol{\Omega} = \sum_{1\le j\le N}\boldsymbol{\omega}_j$. In this way, the dependence of $\boldsymbol{a}_{i\widetilde{g}_i(n)}$ on the strategy valuations is made explicit. Furthermore, one sees that the relevant instantaneous configurational variable of the system is $s_i(n) = \operatorname{sign}\left[y_i(n)\right] \in \{-1, 1\}$. With the above notation, the evolution of the preferences y_i takes the form of the (non-linear) map

$$(2.10)\quad y_i(n+1) = y_i(n) - \frac{1}{N}\,\xi_i^{\mu(n)}\left[A(n) - \eta\; a^{\mu(n)}_{i\widetilde{g}_i(n)}\right] = \\ = y_i(n) - \frac{1}{N}\,\xi_i^{\mu(n)}\left[\Omega^{\mu(n)} + \sum_{1\le j\le N}\xi_j^{\mu(n)} s_j(n) - \eta\;\xi_i^{\mu(n)} s_i(n)\right]$$

2.2. Techniques

The static analysis of this model is based on a mapping of the stationary states of (2.10) onto the minima of a mean-field spin-glass Hamiltonian. The latter are analyzed by the replica method. The direct dynamical study of (2.10) is instead pursued by means of generating functional analysis. In this section we describe briefly the main ideas of these two (by now, standard) techniques.

2.2.1. Replica method. The replica method is the standard way to deal with the equilibrium properties of random systems. In the statistical mechanics framework, typically, a random Hamiltonian $H(\boldsymbol{\sigma})$ is given, which associates to each microscopic spin configuration $\boldsymbol{\sigma} = \{\sigma_i\}_{1\le i\le N}$ a real number (its energy). H is supposed to depend on a family of at least $\mathcal{O}(N)$ random variables per degree of freedom*, which we denote collectively by $\boldsymbol{a}$. A famous example is the Sherrington–Kirkpatrick model [29], defined by $H(\boldsymbol{\sigma}) = (1/\sqrt{N})\sum_{1\le i<j\le N} a_{ij}\sigma_i\sigma_j$, with Ising spins $\sigma_i \in \{-1,1\}$ and Gaussian couplings a_{ij} with zero mean and unitary variance. In this case, $\boldsymbol{a} = \{a_{ij}\}_{1\le i<j\le N}$.

Clearly, one cannot calculate the partition function $Z = \mathrm{Tr}_{\boldsymbol{\sigma}} e^{-\beta H(\boldsymbol{\sigma})}$ for every fixed realization $\boldsymbol{a}$ of the disorder. However, it is to be expected that when $N\to\infty$ intensive quantities like the free-energy density do not depend on the specific $\boldsymbol{a}$ but, rather, only on the statistical properties of the disorder distribution. Therefore one concentrates on such functionals as the disorder-averaged free-energy density, $F = -\lim_{N\to\infty}(\beta N)^{-1}\left[\log Z\right]_{\boldsymbol{a}}$. (Here, and throughout this thesis, $[\;]_{\boldsymbol{a}}$ denotes the average over disorder.) One disposes of the logarithm by the formula

$$(2.11)\qquad \left[\log Z\right]_{\boldsymbol{a}} = \lim_{n\to 0}\frac{\left[Z^n\right]_{\boldsymbol{a}} - 1}{n} \qquad \text{or, equivalently,} \qquad \left[\log Z\right]_{\boldsymbol{a}} = \lim_{n\to 0}\frac{1}{n}\log\left[Z^n\right]_{\boldsymbol{a}}$$

The issue is then to compute $\left[Z^n\right]_{\boldsymbol{a}}$. The replica trick consists in evaluating it for integer n, that is when Z^n can be written as a product of n terms,

$$(2.12)\qquad Z^n = \left(\mathrm{Tr}_{\boldsymbol{\sigma}} e^{-\beta H(\boldsymbol{\sigma})}\right)^n = \prod_{1\le a\le n}\mathrm{Tr}_{\boldsymbol{\sigma}^a} e^{-\beta H(\boldsymbol{\sigma}^a)} = \mathrm{Tr}_{\boldsymbol{\sigma}^1}\cdots\mathrm{Tr}_{\boldsymbol{\sigma}^n} e^{-\beta\sum_{1\le a\le n} H(\boldsymbol{\sigma}^a)}$$

When taking the $n\to 0$ limit, one simply postulates that (2.12) is correct for non-integer values of n as well. Notice that each term $H(\boldsymbol{\sigma}^a)$ contains the same disorder realization $\boldsymbol{a}$ as the original Hamiltonian. It is an exact replica of the original system. Hence, $\sum_{1\le a\le n} H(\boldsymbol{\sigma}^a)$ can be seen as the Hamiltonian of a system of n non-interacting copies of the original system. The computation of the disorder-average of (2.12) is usually straightforward. One finds that the disorder-average generates two-replica quantities like the *overlaps* $Q_{ab} = (1/N)\sum_{1\le i\le N}\sigma_i^a\sigma_i^b$, and the final expression looks like

$$(2.13)\qquad \left[Z^n\right]_{\boldsymbol{a}} = \int e^{-\beta n N f(\mathsf{Q})}\, d\mathsf{Q}$$

where $\mathsf{Q} = \{Q_{ab}\}_{1\le a,b\le n}$ is the overlap matrix and f is some function whose physical meaning will soon become clear. Notice that, as (2.12) is invariant under permutations of the n replicas, so must be f. This implies that f will depend only on algebraic invariants of the group of permutations of n elements, like for instance $\sum_{1\le a,b\le n} Q_{ab}$ or $\sum_{1\le a,b\le n} Q_{ab}^2 = \mathrm{tr}(\mathsf{Q}^2)$.

*The reason for this requirement is that Hamiltonians depending on a smaller number of random variables, e.g. $\mathcal{O}(1)$ per degree of freedom, can be studied without the replica method. The Mattis model, for instance, defined by $H(\boldsymbol{\sigma}) = -(\sum_{1\le i\le N} a_i\sigma_i)^2$, with Ising spins $\sigma_i\in\{-1,1\}$ and random a_i with $\mathrm{Prob}\{a_i = \pm 1\} = 1/2$, can be solved straightforwardly in the new spin variables $\xi_i = a_i\sigma_i$.

When $N \to \infty$, (2.13) can be evaluated by the steepest-descent method. Assuming that $\lim_{N\to\infty}$ and $\lim_{n\to 0}$ commute, so one may take the former first, for the disorder-averaged free energy density F one obtains

$$F = -\lim_{N\to\infty}\frac{1}{\beta N}\left[\log Z\right]_{\boldsymbol{a}} = -\lim_{N\to\infty}\lim_{n\to 0}\frac{1}{\beta n N}\log\left[Z^n\right]_{\boldsymbol{a}} = -\lim_{n\to 0}\lim_{N\to\infty}\frac{1}{\beta n N}\log\left[Z^n\right]_{\boldsymbol{a}} = \simeq -\lim_{n\to 0}\frac{1}{\beta n N}\log e^{-\beta n N\ \mathrm{extr}_{\mathsf{Q}} f(\mathsf{Q})} = \lim_{n\to 0}\mathrm{extr}_{\mathsf{Q}}\ f(\mathsf{Q}) \tag{2.14}$$

One sees that the physical value of the free energy is obtained by evaluating f at a saddle point, where $\partial_{Q_{ab}} f(\mathsf{Q}) = 0$.

The solutions of the saddle-point equations correspond to a point Q^* in the space of $n\times n$ symmetric (since $Q_{ab} = Q_{ba}$) matrices. The selection of the correct saddle points has to be guided by physical intuition. In general, however, one can say that the saddle-point equations will have two distinct types of physical solutions Q^*, depending on whether Q^*_{ab} is independent on the pair (a, b) (in which case the replicas are equivalent and the permutation symmetry is preserved) or not (in which case the replicas are not equivalent, and hence the permutation symmetry is broken). It is possible to show that the first class of solutions describes the ergodic regime. The relevant saddle point has the form

$$\mathsf{Q}^* = q\ \mathsf{E} + (Q - q)\ \mathsf{I} \qquad \text{or, equivalently} \qquad Q^*_{ab} = Q\delta_{ab} + q(1 - \delta_{ab}) \tag{2.15}$$

where E is the $n \times n$ matrix will all entries equal to one and I is the identity matrix. (2.15) is called *replica-symmetric* Ansatz. Typically, it describes correctly the physics in the high-temperature phase.

Solutions Q^* belonging to the second class are said to be *replica-symmetry breaking*. They describe phases with multiple ergodic components. Such solutions can be of many types. The simplest possibility is that

$$\mathsf{Q}^* = q_0\ \mathsf{E}_n + (q_1 - q_0)\ \mathsf{I}_{\frac{n}{m}} \otimes \mathsf{E}_m + (Q - q_1)\ \mathsf{I}_n, \tag{2.16}$$

where subscripts denote the dimensionality of the involved matrices, and $\otimes$ denotes the Kronecker (tensor) product. In this case, Q^*_{ab} is equal to Q on the diagonal, to q_1 on the n/m diagonal blocks of size m, and to q_0 otherwise. (2.16) is called *one-step replica-symmetry breaking* Ansatz. If the system possesses a spin-glass phase, at low temperatures (2.15) becomes unstable against fluctuations that break the permutation symmetry between replicas, and replica-symmetry breaking solutions of the saddle-point equations bifurcate (continuously or not) from the replica-symmetric solution. See [1, 30–34] for some remarkable examples of spin-glasses.

In the minority game, for $\eta = 0$ no replica-symmetry breaking occurs and (2.15) always describes the stationary state [17]. For $\eta > 0$, instead, the replica-symmetric Ansatz becomes unstable at sufficiently low values of α.

2.2.2. Generating functional analysis. The common strategy of all dynamical methods is to derive and solve evolution equations for a small set of suitably chosen macroscopic observables from the (supposedly Markovian) evolution equations of the underlying microscopic system. We employ here the generating functional technique (a.k.a. dynamical mean-field theory, dynamical functional integral, or path-integral formalism) à la De Dominicis. It allows one to perform the disorder average and take the $N \to \infty$ limit exactly. We will describe here the setting for a discrete-time process. The formulation for continuous-time (e.g. Langevin) equations is similar. The starting idea is to consider, rather than the probability of the occurrence of a microscopic state $\boldsymbol{\sigma} = \{\sigma_i\}_{1\le i\le N}$, the probability Prob{path} of finding a microscopic time evolution or path $\boldsymbol{\sigma}(0) \to \boldsymbol{\sigma}(1) \to \cdots \to \boldsymbol{\sigma}(T)$. Using this, one can calculate the moment-generating function

$$Z[\boldsymbol{\psi}] = \left\langle e^{\mathrm{i}\sum_{t\le T}\sum_{1\le i\le N}\psi_i(t)\sigma_i(t)}\right\rangle_{\text{paths}} = \sum_{\text{paths}} e^{\mathrm{i}\sum_{t\le T}\sum_{1\le i\le N}\psi_i(t)\sigma_i(t)}\mathrm{Prob}\{\text{path}\} \tag{2.17}$$

Notice the built-in normalization $Z[\mathbf{0}] = 1$. If $W(\boldsymbol{\sigma}|\boldsymbol{\sigma'})$ denotes the transition probability density for going from $\boldsymbol{\sigma'}$ to $\boldsymbol{\sigma}$ then this is nothing but

$$Z[\boldsymbol{\psi}] = \int e^{\mathrm{i}\sum_{t\le T}\sum_{1\le i\le N}\psi_i(t)\sigma_i(t)} p[\boldsymbol{\sigma}(0)]\prod_{t\le T}\left[W[\boldsymbol{\sigma}(t+1)|\boldsymbol{\sigma}(t)]d\boldsymbol{\sigma}(t)\right] \tag{2.18}$$

where $p[\boldsymbol{\sigma}(0)]$ is the probability distribution of initial conditions.

By taking derivatives of the generating functional with respect to the auxiliary fields $\boldsymbol{\psi} = \{\psi_i\}_{1\leq i\leq N}$ one may obtain all moments of $\boldsymbol{\sigma}$ of all orders at arbitrary times:

$$\big\langle \sigma_{i_1}(t)\sigma_{i_2}(t')\cdots\sigma_{i_k}(t'')\big\rangle_{\rm paths} = \mathrm{i}^{-k}\lim_{\boldsymbol{\psi}\to 0}\frac{\partial^n Z[\boldsymbol{\psi}]}{\partial\psi_{i_1}(t)\partial\psi_{i_2}(t')\cdots\partial\psi_{i_k}(t'')} \tag{2.19}$$

In particular, one is usually interested in certain averages of both one-time quantities like the site-dependent magnetization $m_i(t)$, and two-time quantities like the correlation functions $C_{ij}(t,t')$ and the response functions $G_{ij}(t,t')$. These are given by

$$m_i(t) \equiv \langle\sigma_i(t)\rangle_{\rm paths} = -\mathrm{i}\lim_{\boldsymbol{\psi}\to 0}\frac{\partial Z[\boldsymbol{\psi}]}{\partial\psi_i(t)} \tag{2.20}$$

$$C_{ij}(t,t') \equiv \big\langle\sigma_i(t)\sigma_j(t')\big\rangle_{\rm paths} = -\lim_{\boldsymbol{\psi}\to 0}\frac{\partial^2 Z[\boldsymbol{\psi}]}{\partial\psi_i(t)\partial\psi_j(t')} \tag{2.21}$$

$$G_{ij}(t,t') \equiv \frac{\partial}{\partial\theta_j(t')}\langle\sigma_i(t)\rangle_{\rm paths} = -\mathrm{i}\lim_{\boldsymbol{\psi}\to 0}\frac{\partial^2 Z[\boldsymbol{\psi}]}{\partial\psi_i(t)\partial\theta_j(t')} \tag{2.22}$$

Here, $\boldsymbol{\theta} = \{\theta_i\}_{1\leq i\leq N}$ are external fields which have to be properly added to the dynamics in order to probe the system via perturbations and generate response functions.

This setup is particularly convenient when the microscopic dynamics depends on quenched random variables $\boldsymbol{a}$. In these cases, one can calculate disorder-averaged quantities directly by averaging the generating functional over disorder and taking derivatives. The average over disorder generates two-time quantities like $C_{tt'} = (1/N)\sum_{1\leq i\leq N}\sigma_i(t)\sigma_i(t')$ and $G_{tt'} = -(\mathrm{i}/N)\sum_{1\leq i\leq N}\sigma_i(t)\widehat{\sigma}_i(t')$, where $\widehat{\sigma}_i(t')$ is a conjugate variable of $\sigma_i(t)$, and the final expression looks like

$$[Z[\boldsymbol{\psi}]]_{\boldsymbol{a}} = \int e^{N\phi(\mathsf{C},\mathsf{G})}\,d\mathsf{C}\,d\mathsf{G} \tag{2.23}$$

where $\mathsf{C} = \{C_{tt'}\}$, and similarly for G. Again, one can resort to steepest-descent integration for $N\to\infty$. At the relevant saddle point, it turns out that

$$C_{tt'} = \frac{1}{N}\sum_{1\leq i\leq N}\Big[\big\langle s_i(t)s_i(t')\big\rangle_{\rm paths}\Big]_{\boldsymbol{a}} = -\lim_{\boldsymbol{\psi}\to 0}\frac{1}{N}\sum_{1\leq j\leq N}\frac{\partial^2\,[Z[\boldsymbol{\psi}]]_{\boldsymbol{a}}}{\partial\psi_j(t)\partial\psi_j(t')} \tag{2.24}$$

$$G_{tt'} = \frac{1}{N}\sum_{1\leq i\leq N}\frac{\partial}{\partial\theta_i(t')}\Big[\langle s_i(t)\rangle_{\rm paths}\Big]_{\boldsymbol{a}} = -\lim_{\boldsymbol{\psi}\to 0}\frac{\mathrm{i}}{N}\sum_{1\leq j\leq N}\frac{\partial^2\,[Z[\boldsymbol{\psi}]]_{\boldsymbol{a}}}{\partial\psi_j(t)\partial\theta_j(t')} \tag{2.25}$$

so that when $N\to\infty$ the dynamics is characterized completely by the site- and disorder-averaged correlation and response functions. It usually turns out that the Markovian evolution of the N-spin system is fully described in terms of the non-Markovian dynamics of a single 'effective' spin $\sigma(t)$. The saddle point values of the desired observables have to be solved from the closed set of (saddle-point) equations

$$C_{tt'} = \langle\!\langle\sigma(t)\sigma(t')\rangle\!\rangle \qquad\qquad G_{tt'} = \frac{\partial\,\langle\!\langle\sigma(t)\rangle\!\rangle}{\partial\theta(t')} \tag{2.26}$$

where $\langle\!\langle\ \rangle\!\rangle$ denotes an average over all possible realizations of the non-Markovian stochastic process governing the time-evolution of $\sigma(t)$.

Ergodic stationary states are described by solutions satisfying the following assumptions [4]:

$$\text{Time-translation invariance (TTI):}\qquad \begin{cases}\lim_{t\to\infty}C_{t+\tau,t} = C(\tau)\\ \lim_{t\to\infty}G_{t+\tau,t} = G(\tau)\end{cases} \tag{2.27}$$

$$\text{Finite integrated response (FIR):}\qquad \lim_{t\to\infty}\sum_{t'\leq t}G_{tt'} = \chi < \infty \tag{2.28}$$

$$\text{Weak long-term memory (WLTM):}\qquad \lim_{t\to\infty}\sum_{t'\leq t_w}G_{tt'} = 0\quad\text{for any } t_w \text{ finite} \tag{2.29}$$

In this regime, the correlation and response functions are related by the usual fluctuation-dissipation theorem (FDT) $G(\tau) = -\beta\Theta(\tau)\partial_\tau C(\tau)$.

When ergodicity is broken, any one of the above assumptions might collapse. One possibility, which is realized for instance in the Ising ferromagnet near the Curie-Weiss transition, is that the local susceptibility χ diverges, i.e. FIR fails. In mean-field spin-glasses the collapsing assumption is TTI, whereas both FIR *and* WLTM continue to hold (see [8, 35–37] for some examples of dynamical studies of spin-glasses and [38, 39] for a remarkable connection between statics and dynamics). The fact that time-translation invariance ceases to be valid is reflected in the appearance of solutions (for long times) of the form

$$C_{tt'} = \widehat{C}\left(\frac{h(t')}{h(t)}\right) \qquad \lim_{t,t'\to\infty} \frac{h(t')}{h(t)} < \infty \tag{2.30}$$

where h is an increasing function expressing time reparametrization. However, also in this regime it is possible to obtain analytic results. The reason is the persistence of WLTM. If the system acquired long-term memory, solving the stationary state in the non-ergodic regime would be a hopeless task, since in order to know the asymptotic value of correlation and response functions we would need to know them at all times.

We will see that breakdown of WLTM is precisely what characterizes the minority game.

- For $\eta = 0$, the transition from the ergodic to the non-ergodic regime is signalled by a breakdown of FIR, accompanied by a breakdown of WLTM. TTI is never violated (i.e. there is no aging). This is already bad. But, at least, it is reasonable. A diverging susceptibility should be accompanied by a failure of WLTM (see the definitions above).
- For $\eta > 0$ the situation is yet more complicated. As before, aging is absent, hence TTI always holds. But the system acquires memory. Surprisingly, however, FIR is never violated. So in this case there is long-term memory *and* a finite local susceptibility.

The breakdown of WLTM in the minority game makes it very hard to obtain an analytic solution in the non-ergodic regime. Indeed, we have not been able to derive equations for the persistent order parameters in this case.

At this point, it is important to stress that the fact that spin-glasses don't equilibrate does not mean that there are no stable states. Rather, it is related to the particular initial condition (a random spin configuration) of the dynamics. Due to the high dimensionality of the phase space, such an initial condition will almost surely lie on the boundary between the basins of attraction of some stable states. As a consequence, such states cannot be reached, and the system remains on the boundary forever. This situation is by no means special. A generic ferromagnetic spin-system at low-temperatures has two distinct equilibrium states. But if the relaxation starts from a random (paramagnetic) initial configuration, an infinite system will remain paramagnetic indefinitely. Random configurations are for spin-glasses the most natural initial condition, because, unlike ferromagnetic Hamiltonians, spin-glass Hamiltonians provide no information a priori about the pattern of symmetry breaking. In general, in spin-glasses as in any other system, the dynamical behaviour depends on initial conditions.

2.3. Survey of the $\eta = 0$ case

2.3.1. Numerical results. In the original setup with $\eta = 0$ [13], the information supplied to agents upon which they based their trading decisions consisted of the true history of the market, namely of the past M minority actions (the minority action at round n is just $-\mathrm{sign}[A(n)]$). Numerical simulations, performed with initial conditions of (2.10) given by $y_i(0) = 0$ for all i, revealed a non-trivial dependence of the re-scaled stationary volatility σ_0^2/N on the ratio α of the number 2^M of different information patterns to the number N of agents (see Fig. 2.1 (a), upper panel). For large α, the agents exhibit essentially random behaviour, and the volatility attains the value corresponding to randomly-trading agents, i.e. $\sigma_0^2/N = 1$. As α is reduced, the volatility is found to decrease below the random value, indicating that the system is behaving more efficiently. This implies the onset of a *cooperative* phase, in which the individualistic agents have 'learned' to improve their strategy selection, somehow managing to reduce the global fluctuations below the level corresponding to random traders. The volatility attains a minimum at a value of α given by $\alpha_c \simeq 0.34$. A further decrease of α causes

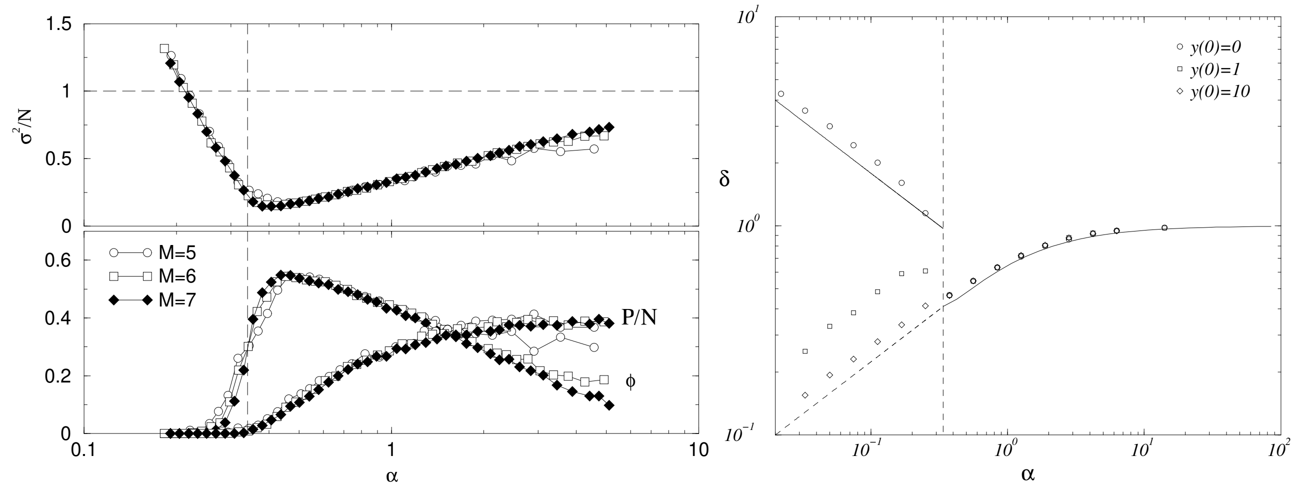

(a) From Ref. [16]

(b) From Ref. [22]

FIGURE 2.1. (a) Upper panel: behaviour of the re-scaled volatility σ_0^2/N as a function of the control parameter $\alpha = 2^M/N$ for different values of the "memory" M, obtained from numerical simulations with initial conditions $y_i(0) = 0$ for all i. As α becomes larger than the displayed values, $\sigma_0^2/N \to 1$. Lower panel: behaviour of the re-scaled predictability P/N and of the fraction ϕ of 'frozen' agents. The vertical dashed line is placed at $\alpha = 0.34 \simeq \alpha_c$. Results are shown for $S = 2$. (b) Behaviour of the square-root of the re-scaled volatility, $\delta = \sigma_0/\sqrt{N}$, as a function of α for different initial conditions $y_i(0) \equiv y(0)$ (all i). The vertical dashed line is at $\alpha = \alpha_c$. The solid and dashed lines in the region $\alpha < \alpha_c$ represent the $\alpha^{-1/2}$ and $\alpha^{1/2}$ behaviours, respectively. The solid line for $\alpha > \alpha_c$ is an analytic estimate derived in [22].

a drastic increase of volatility, and eventually it exceeds the random level, thereby signalling a highly inefficient system, for low values of α.

The existence of a phase transition reminiscent of those found in spin systems was demonstrated by studying numerically the *predictability* P, defined as

$$P = \frac{1}{\alpha N} \sum_{\mu \leq \alpha N} \langle A|\mu\rangle_{\rm time}^2 \qquad \langle A|\mu\rangle_{\rm time} = \lim_{L\to\infty} \frac{1}{L-\ell} \sum_{\substack{\ell \leq n \leq L \\ \mu(n)=\mu}} A(n) \tag{2.31}$$

where $\langle\ |\mu\rangle_{\rm time}$ stands for the time average conditioned on $\mu(n) = \mu$. The predictability (also called *exploitable information*) measures, loosely speaking, the extent to which the winning action is predictable on the basis of the received information pattern. To see this, notice that $P = 0$ when $\langle A|\mu\rangle_{\rm time} = 0$ for all μ. In this case, for no information pattern μ is there an *a priori* preferred minority action. When $P > 0$, instead, there are some μ's for which $\langle A|\mu\rangle_{\rm time}^2 > 0$. For these particular patterns, $\langle A|\mu\rangle_{\rm time}$ is non-zero, so that there is an action that is more likely to be the minority decision, i.e. $-{\rm sign}(\langle A|\mu\rangle_{\rm time})$. In the former case one speaks of an *unpredictable* market, while when $P > 0$ the market is said to be *predictable*. Notice that, like σ^2, P is extensive and so one is mostly concerned with its density P/N. Numerical simulations show that P/N behaves like a physical order parameter (see Fig. 2.1 (a), lower panel) as a function of α, with a singularity at α_c. As more and more agents are present in the system, the market appears to be less and less predictable.

Another interesting quantity that has been measured is the fraction ϕ of "frozen" agents, namely agents whose spin variable $s_i(n) = {\rm sign}[y_i(n)]$ has become either 1 or -1 in the long run. This obviously happens when their preference y_i has become very large, either positive or negative. In particular, one finds that $y_i(n) \sim y_i^* t$ for $t \to \infty$ [15, 17]. Frozen agents always employ the same strategy. ϕ is found to be finite for all $\alpha > \alpha_c$, and it drops steeply to zero below α_c, as shown in Fig. 2.1 (a), lower panel.

The most striking effect observed in numerical simulations is however the lack of ergodicity in the sub-critical phase (see Fig. 2.1 (b)). While for $\alpha > \alpha_c$ the dynamics appears to be ergodic, below α_c the stationary state depends on the initial condition of the dynamics. For $y_i(0) = 0$ (all i) one observes a volatility diverging as $\alpha \to 0$, with $\sigma_0^2/N = \mathcal{O}(\alpha^{-1})$, as also shown in Fig. 2.1 (a). But for different initial conditions the stationary volatility may go to zero with decreasing α. For instance, large initial preferences lead to $\sigma_0^2/N = \mathcal{O}(\alpha)$ as $\alpha \to 0$. The behaviour of ϕ is also found to depend strongly on initial conditions below α_c. Both steady states with a vanishing fraction of frozen agents and ones where a finite fraction (even 1) of agents remains frozen may occur.

The main numerical findings may thus be summarized as follows. When α is sufficiently large, the stationary re-scaled volatility σ_0^2/N approaches the value corresponding to random trading. Reducing α, σ_0^2/N falls below the random value, implying an increase in global efficiency. The dynamic remains ergodic until $\alpha > \alpha_c \simeq 0.34$, and there is a unique stationary state. In all this region, the market is predictable ($P > 0$) and a finite fraction ϕ of the agents is "frozen". At α_c the system apparently undergoes a phase transition to a highly non-ergodic regime ($\alpha < \alpha_c$), where the stationary volatility depends on the initial conditions of the dynamics (i.e. there are multiple stationary states). In particular, both a high- and a low-volatility steady state may occur. A similar behaviour characterizes the fraction of frozen agents. Finally, below α_c the market becomes unpredictable ($P = 0$).

2.3.2. Analytic theory: statics. Attempts at providing a theoretical explanation of the remarkably rich phenomenology that emerged in computer experiments naturally started from the observation that the relevant macroscopic observables, specifically the volatility σ_0^2, can be scaled with N in such a way that they become N-independent in the large N limit. This suggested that a statistical mechanics approach to the model solution could have been successful [16]. The $\eta = 0$ dynamics,

$$y_i(n+1) = y_i(n) - \frac{1}{N}\xi_i^{\mu(n)}\Omega^{\mu(n)} - \frac{1}{N}\sum_{1\leq j\leq N}\xi_i^{\mu(n)}\xi_j^{\mu(n)}s_j(n) \tag{2.32}$$

presented however several difficulties. The first was that, in the original setup where the information patterns presented to agents are the true market histories, it is non-Markovian (actually, it depends on the past M rounds). Secondly, it is non-linear, and the non-linear term (sign$[y_i(n)]$) is singular. As a consequence, it is not possible to re-cast it in the form of a gradient-descent process with a well-defined Lyapunov function, whose minima would then correspond to the desired stationary states. In turn, this implies that (2.10) does not satisfy detailed balance, so that strictly speaking there is no Gibbs-like (thermodynamic) equilibrium. In the light of this fact, the applicability of statistical mechanics' methods is itself questionable.

The first obstacle (non-Markovianness) was overrun when numerical evidence was presented that the stationary state of (2.10) remains (roughly) unchanged if random information, in the form of a randomly chosen integer $\mu \leq \alpha N$, is supplied to agents instead of the real market history [40]. (This is seen to be strictly true only for $\alpha < \alpha_c$, but it holds qualitatively also for $\alpha > \alpha_c$ [41].) The essential is that all agents receive the same information. Reducing (2.32) to a Markovian process, this result allowed for a major simplification.

In order to circumvent the second problem (singularity of the non-linear term), the first attempts at developing a solvable statistical theory were based on a regularization of 2.32. The idea [15] is to introduce an "inverse-temperature" $\Gamma \geq 0$ (to be later interpreted as the *learning rate* of agents) to smoothen the choice mechanism (5.9) to

$$\text{Prob}\{\widetilde{g}_i(n) = g\} = C\ e^{\Gamma p_{ig}(n)} \tag{2.33}$$

with C a normalization constant. Upon varying Γ, one interpolates between random strategy selection, corresponding to $\Gamma = 0$, and the maximum-valuation rule (5.9), which is recovered in the limit $\Gamma \to \infty$. This procedure is essentially equivalent to the introduction of an additive noise in the agents' behaviour, for example by substituting sign$[y_i(n)]$ with sign$[y_i(n) + z_i(n)/\Gamma]$, where $z_i(n)$ are independent identically-distributed zero-average random numbers with unit variance. The advantage of doing so is that it allows the replacement of all the annoying Ising spins $s_i(n) = \text{sign}[y_i(n)] \in \{-1, 1\}$ with the better-looking "soft" spins $\phi_i(n) \equiv \langle s_i(n)\rangle = \tanh[\Gamma y_i(n)] \in [-1,1]$, where $\langle\ \rangle$ denotes a statistical average with probability distribution (2.33). Of course, this is possible at the expense of passing from a system with a discrete configuration space to one with a continuous configuration space. (Surprisingly, the noisy behaviour was found to effect the volatility for $\alpha < \alpha_c$. In particular, it was numerically shown that σ_0^2 increases with Γ (low-volatility states could be reached with a sufficiently small Γ), indicating that the stationary state for $\alpha < \alpha_c$ also depends on Γ.)

In this modified setup, denoting by a bar the average over information patterns (i.e. $\overline{x^\mu} = \sum_{\mu\le\alpha N} x^\mu/(\alpha N)$), the average (long-time) magnitude of fluctuations is easily seen to be given by

$$\sigma^2 \equiv \overline{\langle A^2\rangle} = \sum_{1\le i,j\le N} J_{ij}\phi_i\phi_j + \sum_{1\le i\le N} h_i\phi_i + \sum_{1\le i\le N} J_{ii}(1-\phi_i^2) + \overline{(\Omega^\mu)^2} \tag{2.34}$$

where $\alpha N J_{ij} = \boldsymbol{\xi}_i\cdot\boldsymbol{\xi}_j$ and $\alpha N h_i = 2\boldsymbol{\Omega}\cdot\boldsymbol{\xi}_i$. The stationary volatility (2.6) can be obtained by evaluating (2.34) in the steady state, i.e. with the stationary values of the ϕ_i's. The latter have to be somehow deduced from the "regularized" version of (2.32). After some debate [15,17,20,21] about the continuous-time limit of (2.32), it was finally shown [19,25] that the steady states are correctly described by the minima of the random function

$$H_0(\boldsymbol{\phi}) = \overline{\langle A\rangle^2} = \sum_{1\le i,j\le N} J_{ij}\phi_i\phi_j + \sum_{1\le i\le N} h_i\phi_i + \overline{(\Omega^\mu)^2} \tag{2.35}$$

in $[-1,1]^N$. One can understand this arguing as follows. Consider the average of (2.32) in the stationary state of the modified dynamics. For $N\to\infty$ it is simple to show that it is given by

$$\partial_t \langle y_i\rangle(t) = -\sum_{1\le j\le N} J_{ij}\phi_j - h_i = -\partial_{\phi_i} H_0(\boldsymbol{\phi}) \tag{2.36}$$

where t is a re-scaled time. Motivated by numerical simulations, one looks for solutions of the form $y_i(t) = v_i t$, with constant v_i. If $v_i = 0$, then the stationary value of the preference $\langle y_i\rangle$ is finite, y_i^*, and in the steady state $\phi_i = \tanh[\Gamma y_i^*] \in]-1,1[$. These agents are fickle, they are flipping from one strategy to the other. If $v_i \neq 0$, instead, then $y_i \to \pm\infty$ and $\phi = \text{sign}[v_i] \in \{-1,1\}$. These agents are frozen, in that they employ just one strategy in the long run. One sees that both fickle and frozen agents are described by the minima of H_0. This argument can be demonstrated to be exact for $\Gamma = 0$. In fact, for $\Gamma = 0$, H_0 is a Lyapunov function of the dynamics. This is not true in general for $\Gamma > 0$, but it still provides a very accurate description of the stationary states.

Comparing (2.31) and (2.35), one recognizes that the minimized quantity in this case is actually the predictability. Its minima can be studied by the replica method (a more detailed account will be given for the case of general η, so we omit all details). Within the (correct) replica-symmetric Ansatz, one finds that $\min_{\boldsymbol{\phi}} H_0(\boldsymbol{\phi}) > 0$ for $\alpha > \alpha_c \simeq 0.3374$, while $\min_{\boldsymbol{\phi}} H_0(\boldsymbol{\phi}) = 0$ for $\alpha < \alpha_c$, in full agreement with the numerical results for the predictability (Fig. 2.1 (a), lower panel). The corresponding stationary volatility σ_0^2/N is found to agree with simulations for $\alpha > \alpha_c$ (see Fig. 2.2).

The deviation of the replica prediction from numerical results for $\alpha < \alpha_c$ is obviously related to the observed dependence of the stationary state on the initial conditions. This represents the limit of the statistical-mechanics approach. In fact, as said before, the dynamics (2.32) does not admit a Lyapunov function, so there is no reason why H_0 should describe the correct stationary states of the process at all. However, it is possible to show (but the analysis of this point is quite technical [25], and we won't report it here) that the static approximation provides the correct result in both the ergodic and the non-ergodic regime for $\Gamma = 0$. In this case, in fact, H_0 is exactly a Lyapunov function of the dynamics (2.32) and, remarkably, the replica-symmetric theory also accounts for different initial conditions. For $\Gamma > 0$ and $\alpha < \alpha_c$ a more complex analysis is needed.

Let us finally say that in order to obtain the globally optimal state of the system, one has to study the minima of the magnitude of market fluctuations (2.34) directly [17]. Results are shown in Fig. 2.2. The dynamics having these minima as stationary states can be seen to be given by (2.10) with

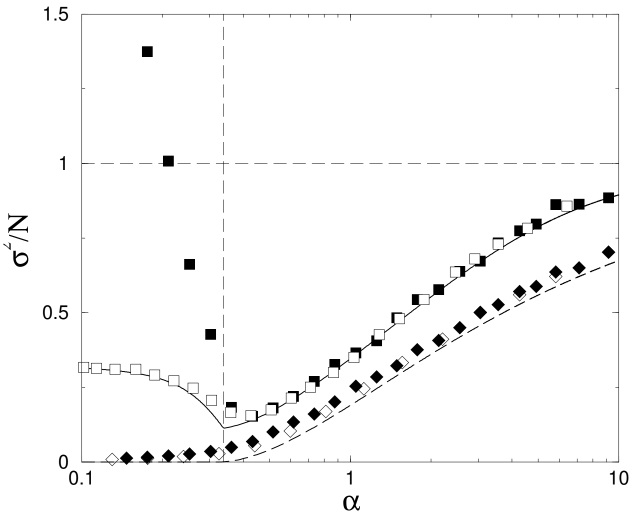

FIGURE 2.2. Upper curves: volatility σ_0^2/N of the $\eta = 0$ model as deducted from numerical simulations (full squares), from numerical minimization of H_0 (open squares), and from the analytic result obtained with the replica method (continuous line). Lower curves: stationary volatility σ_1^2 obtained from the dynamics 2.37. Results are shown from numerical simulations (full diamonds), from numerical minimization of $\sigma^2(n)$ (open diamonds), and from the theoretical lower bound obtained with the replica method within the (incorrect) replica-symmetric approximation. From [17].

$\eta = 1$, namely

$$y_i(n+1) = y_i(n) - \frac{1}{N}\xi_i^{\mu(n)}\Omega^{\mu(n)} - \frac{1}{N}\sum_{1\leq j\leq N}\xi_i^{\mu(n)}\xi_j^{\mu(n)}s_j(n) + [\xi_i^{\mu(n)}]^2 s_i(n) \tag{2.37}$$

So one sees that when agents fully take their market-impact into account, the system's global efficiency improves substantially both above and below the transition. With a modest amount of foresight, one can argue that varying η between 0 and 1 means passing from the upper curve to the lower curve in Fig. 2.2. This will be seen however to happen in a totally non-trivial manner.

2.3.3. Analytic theory: dynamics. Dynamical studies, based on the generating functional technique, initially focused on a 'batch' version of the $\eta = 0$ minority game [22, 23], where one studies the average of (2.10) over information patterns. We defer the reader to Chapter 4 for motivation and details about this model. Let it suffice to say here that the 'batch' minority game is in discrete time (therefore one does not need to construct a proper continuous-time limit), and it yields results which are qualitatively very similar, if not identical, to those of the standard dynamics. For $\eta = 0$ its equations are

$$y_i(t+1) = y_i(t) - \sum_{1\leq j\leq N} J_{ij}s_j(t) - h_i \qquad s_i(t) = \text{sign}[y_i(t)] \tag{2.38}$$

with $J_{ij} = (2/N)\sum_{1\leq\mu\leq\alpha N}\xi_i^\mu\xi_j^\mu$ and $h_i = (2/\sqrt{N})\sum_{1\leq\mu\leq\alpha N}\Omega^\mu\xi_i^\mu$. One finds that for large values of α the dynamics is ergodic and the stationary state can be solved within the assumption of absence of anomalous response (also called finite static susceptibility). This assumption is found to break down at $\alpha_c \simeq 0.3374$, where correspondingly the susceptibility diverges signaling a phase transition, in agreement with the static results. The stationary volatility has been estimated in this region, and the obtained result displays a good agreement with numerical simulations (see Fig. 2.1 (b)). Below α_c, the dynamics is non-ergodic and the stationary state is found to depend strictly on the initial conditions $\{y_i(0)\}$. The square-root of the volatility is found (correctly) to diverge with decreasing α as $\mathcal{O}(\alpha^{-1/2})$ for weak initial preferences, i.e. for all $y_i(0)$ sufficiently close to 0. A critical value $y_c(0)$ was identified, such that for initial strategy valuations above this value the high-volatility solution disappears and is replaced by one with volatility decreasing with α as $\mathcal{O}(\alpha^{1/2})$ (see again Fig. 2.1 (b)). In particular, it turns out that $y_c(0) = (2\pi e)^{-1/2} \simeq 0.242$.

It is worth mentioning that a dynamical solution of the standard non-batch case, with the dynamics given by (2.32), has been recently obtained using a delicate but exact procedure for deriving a continuous-time master equation. The resulting stochastic equations have been again solved with generating functional technique [24].

CHAPTER 3

Statics

We will expose in this chapter the static approximation for the analysis of the stationary states of the dynamics of the minority game (2.10). The problem will be mapped onto that of studying the minima of a mean-field spin-glass Hamiltonian. Ergodic stationary states will be characterized within the replica-symmetric Ansatz (2.15). This will be shown to break down at any $\eta > 0$ for sufficiently small values of α. The AT-line will be evaluated, and it will be shown that in the non-ergodic regime the number of stationary states is exponentially large in the system size. Finally, stationary states in the non-ergodic regime will be characterized upon assuming a one-step replica-symmetry breaking Ansatz. This will provide us with a very good approximation of the true volatility.

3.1. Mapping to a minimization problem

3.1.1. General considerations. We have seen that the instantaneous configuration of the system is conveniently described in terms of the variables $s_i(n) = \text{sign}[y_i(n)] \in \{-1, 1\}$, where the preferences $y_i(n)$ evolve according to (2.10), that is

$$y_i(n+1) = y_i(n) - \frac{1}{N}\,\xi_i^{\mu(n)}\Omega^{\mu(n)} - \frac{1}{N}\sum_{1\leq j\leq N}\xi_i^{\mu(n)}\xi_j^{\mu(n)}s_j(n) + \frac{\eta}{N}\,[\xi_i^{\mu(n)}]^2 s_i(n) \tag{3.1}$$

The idea behind the static approach to the minority game is to find a Lyapunov function (i.e. a function of the dynamic state variables that decreases monotonically along the system's trajectories and is bounded from below), and use equilibrium techniques of statistical mechanics, like the replica method, to study its minima, which would correspond to the stationary states of the dynamics. This program only applies when the system reaches a stable stationary state (and, of course, when $N \to \infty$). Therefore, it admits no Lyapunov function. However, one can try to circumvent this problem by regularizing the microscopic dynamics with some additive noise (or inverse temperature Γ), as done with the original model. Following either [17–19] or the lines traced in the previous chapter, it is not difficult to understand that, for general η, the regularized dynamics (approximately) minimizes the random function

$$H_\eta(\boldsymbol{\phi}) = \sum_{1\leq i,j\leq N} J_{ij}\phi_i\phi_j + \sum_{1\leq i\leq N} h_i\phi_i + \eta\sum_{1\leq i\leq N} J_{ii}[1-\phi_i^2] + \overline{(\Omega^\mu)^2} \tag{3.2}$$

with $\alpha N J_{ij} = \boldsymbol{\xi}_i\cdot\,\boldsymbol{\xi}_j$, $\alpha N h_i = 2\boldsymbol{\Omega}\cdot\boldsymbol{\xi}_i$, and $\alpha N\overline{(\Omega^\mu)^2} = \boldsymbol{\Omega}\cdot\boldsymbol{\Omega}$. Using (2.35) this can also be written

$$H_\eta(\boldsymbol{\phi}) = \overline{\langle A\rangle^2} + \eta\sum_{1\leq i\leq N}\overline{(\xi_i^\mu)^2}(1-\phi_i^2) \tag{3.3}$$

where $A^\mu = \Omega^\mu + \sum_{1\leq i\leq N}\xi_i^\mu s_i(n)$ is the total bid defined in (2.9). In the special cases $\eta = 0$ we have seen that H_0 is exactly the predictability P. For $\eta = 1$ one finds $H_1(\boldsymbol{\phi}) = \overline{\langle A^2\rangle}$, and the dynamics minimizes the volatility (see (2.34)).

A full characterization of the minima of H_η requires the replica method, and is the subject of this chapter. However, to have an idea of how the introduction of a non-zero η in the dynamics affects the stationary volatility, we can anticipate the theoretical analysis that will follow, and show the results obtained from the replica method within the replica-symmetric Ansatz for different values of η (see Fig. 3.1). The reader is warned that the replica-symmetric Ansatz will be seen to provide an exact solution for $\eta \leq 0$, while for $\eta > 0$ it is only valid for sufficiently large values of α. We will discuss this issue in detail. However, this approximation is enough for our present purpose, since even in the worst case ($\eta = 1$, where it fails for all α) it provides a fair lower bound to the true volatility. One

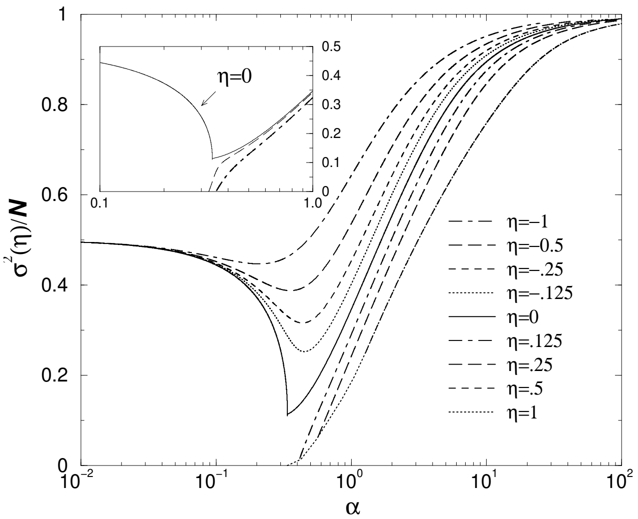

FIGURE 3.1. Theoretical estimate of the volatility σ_η^2/N as a function of α for different values of η (curves from top to bottom are for increasing values of η) from -1 to 1. Inset: blow-up of the situation for small positive values of η.

sees that the volatility continuously decreases as η increases for $\alpha > \alpha_c$, whereas for low values of α it has a jump discontinuity from a finite value at $\eta \leq 0$ to zero as soon as $\eta > 0$. This shows that as soon as agents starts considering their market-impact, i.e. for an infinitesimally small $\eta > 0$, the global efficiency is improved with respect to the $\eta = 0$ case; in particular for sufficiently low α (i.e. when many agents participate the market) fluctuations are reduced by a finite amount.

Finally, we report of an important property of H_1, which will be used several times in the following is that it attains its minima in the corners of the configurations space $[-1,1]^N$, hence on $\{-1,1\}^N$. This follows from harmonicity, i.e. from the fact that $\nabla^2_{\boldsymbol{\phi}} H_1(\boldsymbol{\phi}) = 0$ (one may verify by a straightforward calculation that $\partial^2_{\phi_i} H_1(\boldsymbol{\phi}) = 0$ for all i). This implies that when $\eta = 1$ in the stationary state all agents are frozen.

3.1.2. Dynamical instability for $\Gamma = 0$. There exists a simple and elegant argument developed by Marsili to show that for $\Gamma = 0$ there is an instability in the system's dynamics for $\eta > 0$. Starting from (3.1), and averaging over the stationary state one obtains

$$\partial_t y_i(t) = -\sum_{1\leq j\leq N} J_{ij}\phi_j - h_i + \eta J_{ii}\phi_i \tag{3.4}$$

Let us argue as done for the $\eta = 0$ case. When $t \to \infty$ will have a group of agents for which $|y_i|$ grows with time, e.g. $y_i(t) = v_i t$. We call these agents *frozen*, because they always employ the same strategy. For them, $\phi_i = \text{sign}(v_i)$. On the other hand, there will be *fickle* agents, for which $\phi_i = \tanh(y_i^*)$, where $y_i^* = \lim_{t\to\infty} y_i(t)$. Let ϕ denote the fraction of frozen agents and $\overline{\phi} = 1-\phi$ denote the fraction of fickle ones, and let us consider the dynamics of fickle agents. Setting $y_i(t) = y_i^* + \epsilon_i(t)$ for these, and expanding (3.4) in powers of $\epsilon_i(t)$ up to the first order one finds

$$\partial_t \epsilon_i(t) = -\sum_{j\leq\overline{\phi}N} \overline{\xi_i^\mu \xi_j^\mu}(1-\phi_j^2)\epsilon_j(t) + \eta\overline{(\xi_i^\mu)^2}(1-\phi_i^2)\epsilon_i(t) = -\sum_{j\leq\overline{\phi}N} T_{ij}\epsilon_j(t) \tag{3.5}$$

where $T_{ij} = \overline{\xi_i^\mu \xi_j^\mu}(1-\phi_j^2) - \eta\overline{(\xi_i^\mu)^2}(1-\phi_i^2)\delta_{ij}$. As long as the matrix $\mathsf{T} = \{T_{ij}\}_{i,j\leq\overline{\phi}N}$ is positive definite, the dynamical system (3.5) is (linearly) stable. Notice that $\mathsf{T} = \mathsf{UV}$ with $U_{ij} = \overline{\xi_i^\mu \xi_j^\mu} - \eta\overline{(\xi_i^\mu)^2}\delta_{ij}$ and $V_{ij} = (1-\phi_i^2)\delta_{ij}$, so that, since all eigenvalues of V are positive definite (because $\phi_i^2 < 1$ for fickle agents), $\det\mathsf{T}$ vanishes when $\det\mathsf{U}$ vanishes. The spectrum of U can be evaluated using the results of [42]. In particular, for our purposes it suffices to calculate the minimum eigenvalue, which we find to be $\lambda_{\min} = (1/2)[(1-\sqrt{\overline{\phi}/\alpha})^2 - \eta]$. Its vanishing signals the onset of instability. This occurs when

$$\overline{\phi} = \alpha\,(1-\sqrt{\eta})^2 \tag{3.6}$$

We have derived this condition in the very special case $\Gamma = 0$. It is remarkable that not only will this turn out to be the AT-line of our model in the static solution, but also the line where long-term memory sets in in the dynamical solution (in which $\Gamma = \infty$ from the outset).

3.2. Solution of the minimization problem by the replica method

In order to solve the minimization problem for H_η one can resort to statistical mechanics techniques, because

$$\lim_{N\to\infty} \min_{\boldsymbol{\phi}\in[-1,1]^N} \frac{H_\eta(\boldsymbol{\phi})}{N} = -\lim_{\beta\to\infty}\lim_{N\to\infty}\frac{1}{\beta N}\left[\log Z\right]_{\boldsymbol{a}} \tag{3.7}$$

where $Z = \int e^{-\beta H_\eta(\boldsymbol{\phi})}\,d\boldsymbol{\phi} \equiv \mathrm{Tr}_{\boldsymbol{\phi}} e^{-\beta H_\eta(\boldsymbol{\phi})}$ is the canonical partition function associated to H_η, and $[\,]_{\boldsymbol{a}}$ denotes the average over disorder, i.e. over all possible choices of the strategies $\{\boldsymbol{a}_{ig}\}$, which are assumed to have the probability distribution, $P(a^\mu_{ig}) = \frac{1}{2}(\delta_{a^\mu_{ig},+1} + \delta_{a^\mu_{ig},-1})$, $g = 1,2$. The calculation of $[\log Z]_{\boldsymbol{a}}$ requires the replica formalism. We hence use the fact that

$$\left[\log Z\right]_{\boldsymbol{a}} = \lim_{n\to 0}\frac{1}{n}\log\left[Z^n\right]_{\boldsymbol{a}} \tag{3.8}$$

and concentrate on $[Z^n]_{\boldsymbol{a}}$.

Letting $\mathcal{A}^a = \Omega^\mu + \sum_{1\le i\le N}\xi^\mu_i\phi^a_i$, we can write Z^n, as

$$\begin{aligned} Z^n &= \mathrm{Tr}_{\boldsymbol{\phi}^1}\cdots\mathrm{Tr}_{\boldsymbol{\phi}^n}\, e^{-\frac{\beta}{\alpha N}\sum_{1\le a\le n}\sum_{\mu\le\alpha N}\left\{(\mathcal{A}^a)^2 + \eta\sum_{1\le i\le N}(\xi^\mu_i)^2\left[1-(\phi^a_i)^2\right]\right\}} \\ &= \mathrm{Tr}_{\boldsymbol{\phi}^1}\cdots\mathrm{Tr}_{\boldsymbol{\phi}^n}\; E_{\boldsymbol{z}}\left[h(\boldsymbol{z})\right] e^{-\frac{\eta\beta}{\alpha N}\sum_{1\le i\le N}\sum_{\mu\le\alpha N}(\xi^\mu_i)^2\sum_{1\le a\le n}\left[1-(\phi^a_i)^2\right]} \end{aligned} \tag{3.9}$$

$$h(\boldsymbol{z}) = e^{-i\sqrt{\frac{2\beta}{\alpha N}}\sum_{1\le a\le n}\sum_{\mu\le\alpha N} z^\mu_a\mathcal{A}^a} \tag{3.10}$$

where $E_{\boldsymbol{z}}(\cdots)$ denotes an average over the $n\alpha N$ unit Gaussian random variables z^μ_a ($a\le n$; $\mu\le\alpha N$) we have introduced in order to linearize the exponent:

$$E_{\boldsymbol{z}}(\cdots) = (2\pi)^{-\frac{n\alpha N}{2}}\int_{\mathbb{R}^{n\alpha N}}\cdots\; e^{-\frac{1}{2}\sum_{1\le a\le n}\sum_{\mu\le\alpha N}(z^\mu_a)^2}\boldsymbol{dz} \tag{3.11}$$

Now we simplify our lives by noticing that, asymptotically, $(\alpha N)^{-1}\sum_{\mu\le\alpha N}(\xi^\mu_i)^2 = 1/2$ for all i, so that

$$\left[Z^n\right]_{\boldsymbol{a}} = \mathrm{Tr}_{\boldsymbol{\phi}^1}\cdots\mathrm{Tr}_{\boldsymbol{\phi}^n}\; E_{\boldsymbol{z}}\left[\left[h(\boldsymbol{z})\right]_{\boldsymbol{a}}\right] e^{-\frac{\eta\beta}{2}\sum_{1\le i\le N}\sum_{1\le a\le n}\left[1-(\phi^a_i)^2\right]} \tag{3.12}$$

Then, calculating explicitly the disorder average taking into consideration the fact that one can factorize over μ and i, we obtain

$$\left[h(\boldsymbol{z})\right]_{\boldsymbol{a}} = e^{-\frac{\beta}{2\alpha}\sum_{1\le a,b\le n}\sum_{\mu\le\alpha N} z^\mu_a z^\mu_b\left(1+\frac{1}{N}\sum_{1\le i\le N}\phi^a_i\phi^b_i\right)} \tag{3.13}$$

At this point we introduce the *overlaps* $Q_{ab} = (1/N)\sum_{1\le i\le N}\phi^a_i\phi^b_i$ by inserting the identity

$$1 = \int d\mathsf{Q}\prod_{a,b=1}^{n}\delta\left(Q_{ab} - \frac{1}{N}\sum_{1\le i\le N}\phi^a_i\phi^b_i\right) = (2\pi)^{-n^2}\int e^{-i\sum_{a,b=1}^{n}\widehat{Q}_{ab}\left(Q_{ab}-\frac{1}{N}\sum_{1\le i\le N}\phi^a_i\phi^b_i\right)}d\mathsf{Q}\;d\widehat{\mathsf{Q}} \tag{3.14}$$

For each term (a,b) in (3.13) we get a factor

$$\begin{aligned} &\frac{1}{2\pi}\int e^{-i\widehat{Q}_{ab}Q_{ab} - \frac{\beta}{2\alpha}\sum_{\mu\le\alpha N} z^\mu_a z^\mu_b + i\left(\widehat{Q}_{ab} - \frac{i\beta}{2\alpha}\sum_{\mu\le\alpha N} z^\mu_a z^\mu_b\right)\frac{1}{N}\sum_{1\le i\le N}\phi^a_i\phi^b_i}\; dQ_{ab}\; d\widehat{Q}_{ab} = \\ &\qquad = \frac{1}{2\pi}\int e^{-i\overline{Q}_{ab}\left(Q_{ab} - \frac{1}{N}\sum_{1\le i\le N}\phi^a_i\phi^b_i\right) - \frac{\beta}{2\alpha}\sum_{\mu\le\alpha N} z^\mu_a z^\mu_b(1+Q_{ab})}\; dQ_{ab}\; d\overline{Q}_{ab} \end{aligned} \tag{3.15}$$

where $\overline{Q}_{ab} = \widehat{Q}_{ab} - \frac{i\beta}{2\alpha} \sum_{\mu \le \alpha N} z_a^\mu z_b^\mu$. Performing the Gaussian integrals over $\boldsymbol{z}$ we easily obtain

$$(2\pi)^{-\frac{n\alpha N}{2}} \int e^{-\frac{1}{2}\sum_{1\le a\le n}\sum_{\mu\le\alpha N}(z_a^\mu)^2 - \frac{\beta}{2\alpha}\sum_{1\le a,b\le n}\sum_{\mu\le\alpha N} z_a^\mu z_b^\mu (1+Q_{ab})}\, d\boldsymbol{z} = \tag{3.16}$$

$$= \int e^{-\frac{1}{2}\sum_{a,b=1}^n T_{ab}\sum_{\mu\le\alpha N} z_a^\mu z_b^\mu}\, d\boldsymbol{z} = (\det \mathsf{T})^{-\frac{\alpha N}{2}} \equiv e^{-\frac{\alpha N}{2}\log\det\mathsf{T}}$$

where $T_{ab} = \delta_{ab} + \frac{\beta}{\alpha}(1 + Q_{ab})$ or, equivalently,

$$\mathsf{T} = \mathsf{I} + \frac{\beta}{\alpha}\,(\mathsf{E} + \mathsf{Q}) \tag{3.17}$$

where I and E are the n-dimensional unit matrix and the n-dimensional matrix with all entries equal to one, respectively. Neglecting all constant prefactors (which all together tend to one as $n \to 0$), we have

$$[Z^n]_{\boldsymbol{a}} = \int e^{-\frac{\alpha N}{2}\log\det\mathsf{T} - i\sum_{1\le a,b\le n}\overline{Q}_{ab}Q_{ab} - \frac{\eta\beta N}{2}\sum_{1\le a\le n}(1-Q_{aa})} \times \tag{3.18}$$

$$\times\, \mathrm{Tr}_{\boldsymbol{\phi}^1}\cdots\mathrm{Tr}_{\boldsymbol{\phi}^n}\left[e^{\frac{i}{n}\sum_{1\le a,b\le n}\overline{Q}_{ab}\sum_{1\le i\le N}\phi_i^a\phi_i^b}\right]\, d\mathsf{Q}\; d\overline{\mathsf{Q}}$$

Factorizing the argument of the traces over i,

$$\left(\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}}\, e^{\frac{i}{n}\sum_{1\le a,b\le n}\overline{Q}_{ab}\phi^a\phi^b}\right)^N = e^{N\log\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}}\, e^{\frac{i}{n}\sum_{1\le a,b\le n}\overline{Q}_{ab}\phi^a\phi^b}} \tag{3.19}$$

and setting, without elaborate justification, $\overline{Q}_{ab} = -\frac{i\beta^2\alpha N}{2}R_{ab}$, we finally arrive at

$$[Z^n]_{\boldsymbol{a}} = \int e^{-\beta n N f(\mathsf{Q},\mathsf{R})}\, d\mathsf{Q}\; d\mathsf{R} \tag{3.20}$$

with

$$f(\mathsf{Q},\mathsf{R}) = \frac{\alpha}{2\beta n}\,\log\det\mathsf{T} + \frac{\eta}{2n}\sum_{1\le a\le n}(1-Q_{aa}) + \frac{\alpha\beta}{2n}\sum_{1\le a,b\le n} R_{ab}\,Q_{ab} + \tag{3.21}$$

$$-\,\frac{1}{\beta n}\log\left[\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}} e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n}R_{ab}\phi^a\phi^b}\right]$$

(3.21) is our desired result. We can integrate (3.20) using the saddle-point technique, by which it is simple to show that

$$\lim_{N\to\infty}\,\min_{\boldsymbol{\phi}\in[-1,+1]^N}\frac{H_\eta(\boldsymbol{\phi})}{N} = \lim_{\beta\to\infty}\lim_{n\to 0} f(\mathsf{Q}^*,\mathsf{R}^*) \tag{3.22}$$

where $(\mathsf{Q}^*,\mathsf{R}^*)$ is a saddle point of (3.21). As usual, we will start by assuming a replica-symmetric saddle-point. This will provide a description of the ergodic regime.

3.3. Replica-symmetric theory

3.3.1. Replica-symmetric approximation. We saw that ergodicity implies that the relevant saddle-point Q^* has the replica-symmetric form

$$\mathsf{Q} \equiv \mathsf{Q}_{\mathrm{RS}} := q\;\mathsf{E} + (Q-q)\;\mathsf{I} \qquad \text{or} \qquad Q_{ab} = \begin{cases} Q & \text{for } a = b \\ q & \text{otherwise}\end{cases} \tag{3.23}$$

We will now assume this form of Q^*. At the same time, an analogous Ansatz for R^* will be made. With this simple choice we can diagonalize the matrix T, $T_{ab} = [1 + \frac{\beta}{\alpha}(Q-q)]\delta_{ab} + \frac{\beta}{\alpha}(1+q)$, quite straightforwardly, obtaining

$$
\begin{array}{llll}
 & \text{Eigenspace :} & \text{Eigenvalue :} & \text{Degeneracy :} \\
(3.24) & \sum_{1\le a\le n} x_a = 0 & \lambda_1 = 1 + \frac{\beta}{\alpha}(Q-q) & n-1 \\
 & \boldsymbol{x} = (1,\ldots,1) & \lambda_2 = \lambda_1 + n\frac{\beta}{\alpha}(1+q) & 1
\end{array}
$$

For the different terms of the free energy we get, respectively

$$
\lim_{n\to 0}\frac{1}{n}\,\log\det\mathsf{T} = \lim_{n\to 0}\frac{1}{n}\,\log\left(\lambda_1^{n-1}\lambda_2\right) = \log\left[1+\frac{\beta}{\alpha}(Q-q)\right] + \frac{\beta(1+q)}{\alpha+\beta(Q-q)}
$$

$$
\lim_{n\to 0}\frac{1}{n}\sum_{1\le a\le n}(1-q_{aa}) = 1-Q \qquad \lim_{n\to 0}\frac{1}{n}\sum_{1\le a,b\le n} R_{ab}Q_{ab} = RQ - rq
$$

$$
\begin{aligned}
\lim_{n\to 0}\frac{1}{n}\log&\left[\mathrm{Tr}_{\{\phi^1,\ldots,\phi^n\}}e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n}R_{ab}\phi^a\phi^b}\right] = \\
&= \lim_{n\to 0}\frac{1}{n}\log\left[\mathrm{Tr}_{\{\phi^1,\ldots,\phi^n\}}e^{\frac{\alpha\beta^2 r}{2}\left(\sum_{1\le a\le n}\phi^a\right)^2+\frac{\alpha\beta^2(R-r)}{2}\sum_{1\le a\le n}(\phi^a)^2}\right] = \\
&= \lim_{n\to 0}\frac{1}{n}\,\log\,E_z\left[\mathrm{Tr}_{\{\phi^1,\ldots,\phi^n\}}e^{\sqrt{\alpha r}\beta z\sum_{1\le a\le n}\phi^a+\frac{\alpha\beta^2(R-r)}{2}\sum_{1\le a\le n}(\phi^a)^2}\right] = \\
&= \lim_{n\to 0}\frac{1}{n}\,\log\,E_z\left[\left(\int_{-1}^{1}e^{-\beta V_z(\phi)}d\phi\right)^n\right] \simeq \lim_{n\to 0}\frac{1}{n}\,\log\,E_z\left[1+n\log\left(\int_{-1}^{1}e^{-\beta V_z(\phi)}d\phi\right)\right] = \\
&= E_z\left[\log\int_{-1}^{1}e^{-\beta V_z(\phi)}d\phi\right]
\end{aligned}
$$

where $E_z(\cdots)$ denotes an average over the unit Gaussian random variable z (namely $E_z(\cdots) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty}\cdots e^{-\frac{1}{2}z^2}dz$) and

$$
V_z(\phi) = -\sqrt{\alpha r}z\phi - \frac{\alpha\beta}{2}(R-r)\phi^2 \tag{3.25}
$$

The RS 'free energy' finally reads

$$
\begin{aligned}
f(\mathsf{Q}_{\mathrm{RS}},\mathsf{R}_{\mathrm{RS}}) = \frac{\alpha}{2\beta}\log\left[1+\frac{\beta(Q-q)}{\alpha}\right] + \frac{\alpha}{2}\frac{1+q}{\alpha+\beta(Q-q)} + \\
+\frac{\alpha\beta}{2}(RQ-rq) - \frac{1}{\beta}\,E_z\left[\log\int_{-1}^{1}e^{-\beta V_z(\phi)}d\phi\right] + \frac{\eta}{2}(1-Q)
\end{aligned}
\tag{3.26}
$$

The corresponding saddle point equations, obtained by variation of (3.26) with respect to Q, q, R and r, turn out to be given by

$$
r = \frac{1+q}{\alpha^2(1+\chi)^2} \qquad \beta(R-r) = -\frac{1}{\alpha}\left(\frac{1}{1+\chi}-\eta\right) \tag{3.27}
$$

$$
Q = E_z\left(\langle\phi^2\rangle_V\right) \qquad \beta(Q-q) = \frac{1}{\sqrt{\alpha r}}\,E_z\left(\langle z\phi\rangle_V\right) \tag{3.28}
$$

where we have defined

$$
\chi = \frac{\beta}{\alpha}(Q-q) \qquad \text{and} \qquad \langle\cdots\rangle_V := \frac{\int_{-1}^{1}\cdots\; e^{-\beta V_z(\phi)}d\phi}{\int_{-1}^{1}e^{-\beta V_z(\phi)}d\phi} \tag{3.29}
$$

Being interested in the limit $\beta\to\infty$ we may look for solutions with $Q=q$ (or χ finite) and $R=r$. For $\eta=0$ such a solution is characterized by a phase transition at $\alpha_c \simeq 0.3374\ldots$ separating an unpredictable phase ($\alpha<\alpha_c$) with $H_0\equiv P/N=0$ from a predictable one ($\alpha>\alpha_c$) with $H_0>0$. The "spin susceptibility" $\chi=\beta(Q-q)/\alpha$ diverges as $\alpha\to\alpha_c^+$.

Starting from the saddle-point equations, it is possible to obtain a closed set of equations for χ, Q, R and $\beta(R-r)$ (we do not report these equations explicitly here). In Fig. 3.2 we compare simulation results for the stationary value of Q with the replica-symmetric prediction. For $\eta\le 0$ one sees that

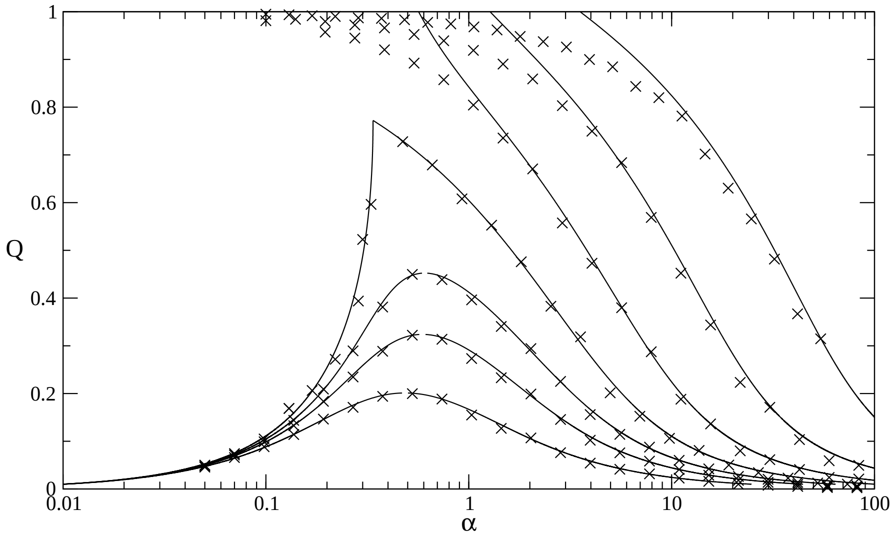

FIGURE 3.2. Stationary value of the self-overlap Q as a function of α for different values of η (from top to bottom, $\eta = -1, -0.5, -0.25, 0, 0.25, 0.5, 0.7$). Solid lines represent the replica-symmetric predictions. Simulations have been run for 500 time steps after equilibration, with $\alpha N = 10^4$. Averages are taken over 100 disorder samples.

the agreement is excellent. This suggests that no ergodicity breaking occurs. For $\eta = 0$ the agreement is equally good. Below $\alpha_c \simeq 0.3374$, we have reported the solution corresponding to an infinitesimal negative η. When $\eta > 0$, instead, the replica-symmetric solution provides the correct result only for sufficiently large values of α. The deviation is, in this case, related to the onset of replica-symmetry breaking.

3.3.2. Failure of the replica-symmetric solution at $\eta = 1$. We will now show that for $\eta = 1$ the RS Ansatz does not provide the correct solution for any α. To see this, it is sufficient to calculate the 'entropy density' corresponding to the replica-symmetric solution with $\eta = 1$, given by $\Sigma = \beta^2 \partial_\beta f$. Recalling that the minima of H_1 occur on the corners of the configuration space $[-1,1]^N$, so that continuous spins ϕ_i become Ising spins σ_i, we have that Σ should be non-negative for all α. Instead, we find that Σ becomes negative in the limit $\beta \to \infty$. We thus conclude that the replica-symmetric solution is incorrect for $\eta = 1$ at all values of α. Since $Q = 1$ when $\eta = 1$, we let $\chi = \beta(1-q)/\alpha$. Moreover, we turn all integrals over $[-1,1]$ to sums over $\{-1,1\}$. Straightforward differentiation of Eq. (3.26) gives

$$\Sigma = \frac{\alpha}{2}\left[\frac{\chi}{1+\chi} - \log(1+\chi)\right] - \frac{\beta\chi(1+q)}{2(1+\chi)^2} + \frac{\alpha^2\beta R\chi}{2} + \\ + \frac{\alpha\beta^2}{2}(R-r) + \left(1 - \beta\frac{\partial}{\partial\beta}\right) E_z\left[\log\sum_{\sigma=\pm1} e^{-\beta V_z(\sigma)}\right] \tag{3.30}$$

The second and third term on the right-hand side cancel exactly, since in the RS solution $R = r = \frac{1+q}{\alpha^2(1+\chi)^2}$ when $\beta \to \infty$. Using Eq. (3.25) one obtains easily

$$E_z\left[\log\sum_{\phi=\pm1} e^{-\beta V_z(\phi)}\right] = \frac{\alpha\beta^2}{2}(R-r) + E_z\left\{\log\left[2\cosh(\beta z\sqrt{\alpha r})\right]\right\} \tag{3.31}$$

from which

$$\Sigma = \frac{\alpha}{2}\left[\frac{\chi}{1+\chi} - \log(1+\chi)\right] + E_z\left\{\log\left[2\cosh(\beta z\sqrt{\alpha r})\right]\right\} - \beta\sqrt{\alpha r}\; E_z\left[z\tanh(\beta z\sqrt{\alpha r})\right] \tag{3.32}$$

follows. Finally, the last two terms on the right-hand side cancel exactly because when $\beta \to \infty$

$$E_z\left[z\tanh(\beta z\sqrt{\alpha r})\right] \to E_z|z| = \sqrt{\frac{2}{\pi}} \qquad E_z\left\{\log\left[2\cosh(\beta z\sqrt{\alpha r})\right]\right\} \to \beta\sqrt{\frac{2\alpha r}{\pi}} \tag{3.33}$$

so that the final result is

$$\lim_{\beta\to\infty}\Sigma = \frac{\alpha}{2}\left[\frac{\chi}{1+\chi} - \log(1+\chi)\right] \tag{3.34}$$

The expression of χ as a function of α may be obtained from the second saddle-point equation in (3.28), that for the case in study ($Q = 1$) takes the form

$$\beta(1-q) = \frac{E_z\left[z\tanh(\beta z\sqrt{\alpha r})\right]}{\sqrt{\alpha r}} \xrightarrow{\beta\to\infty} \alpha\chi = \sqrt{\frac{2}{\alpha\pi r}} \tag{3.35}$$

From this one deduces $r = 2/(\alpha^3\pi\chi^2)$. Comparing this with the saddle point result $r = (1+q)/[\alpha^2(1+\chi)^2] = 2\alpha^{-2}(1+\chi)^{-2}$ one obtains an equation for χ whose only physical (i.e. positive) solution is $\chi = 1/(\sqrt{\alpha\pi}-1)$. With this value of χ, one sees that Eq. (3.34) is negative for all $\alpha > 1/\pi$ and it diverges as $\alpha \downarrow 1/\pi$. For all $\alpha < 1/\pi$ one finds that the zero-temperature entropy density is equal to $-\infty$.

3.4. Stability of the replica-symmetric solution

The RS solution has been just shown to give rise to a negative entropy for all α at $\eta = 1$, indicating its failure since at $\eta = 1$ the stationary state is described by Ising variables. The stability of the RS solution for general η is related to the sign of the Hessian of the replica free energy in the replica-symmetric state. When $\mathcal{H}[f]$ changes sign the RS Ansatz no longer describes the correct extremum of f and a different approximation has to be used. The de Almeida-Thouless (AT) instability refers in fact to the bifurcation of a saddle-point solution without replica symmetry from the replica-symmetric one. The stability matrix is

$$\mathsf{S} = \begin{pmatrix} \mathsf{A} & \mathsf{C} \\ \mathsf{C} & \mathsf{B} \end{pmatrix} \tag{3.36}$$

where the sub-matrices A, B and C have elements

$$A_{(ab)(cd)} = \frac{\partial^2(nf)}{\partial Q_{ab}\partial Q_{cd}} \qquad B_{(ab)(cd)} = \frac{\partial^2(nf)}{\partial R_{ab}\partial R_{cd}} \qquad C_{(ab)(cd)} = \frac{\partial^2(nf)}{\partial Q_{ab}\partial R_{cd}} \tag{3.37}$$

respectively. Here, f is the replica free energy (3.21). We will now calculate the matrix elements of S in the replica-symmetric state. Our ultimate goal is to evaluate the sign of det$\mathsf{S}$.

For A we obtain

$$A_{(ab)(cd)} = -\frac{\beta}{\alpha}\left[\left(\mathsf{T}^{-1}\right)_{ac}\left(\mathsf{T}^{-1}\right)_{bd} + \left(\mathsf{T}^{-1}\right)_{ad}\left(\mathsf{T}^{-1}\right)_{bc}\right] \tag{3.38}$$

When $\mathsf{Q} = \mathsf{Q}_{\mathrm{RS}}$, T takes the simple form $\mathsf{T} = [1+\frac{\beta}{\alpha}(Q-q)]\,\mathsf{I} + \frac{\beta}{\alpha}(1+q)\,\mathsf{E}$ and one may easily verify that $\mathsf{T}^{-1} = (1/\lambda_1)\mathsf{I} - [(\frac{\beta}{\alpha}(1+q))/(\lambda_1\lambda_2)]\mathsf{E}$ where λ_1 and λ_2 are the eigenvalues of T. A is now seen to have three different types of matrix elements in the RS state*, depending on whether none, one or two replica indices of the pair (ab) equal those of the pair (cd). In the $n \to 0$ limit one finds

$$A_{(ab)(ab)} \equiv A_1 = -\frac{\beta}{\alpha}\left[C_1^2 + C_2^2\right] \tag{3.39}$$

$$A_{(ab)(ac)} \equiv A_2 = -\frac{\beta}{\alpha}\,C_1\left[C_1 + C_2\right] \tag{3.40}$$

$$A_{(ab)(cd)} \equiv A_3 = -\frac{2\beta}{\alpha}\,C_1^2 \tag{3.41}$$

where $C_1 = -\frac{\beta(1+q)}{\alpha(1+\chi)^2}$ and $C_2 = \frac{1}{1+\chi} + C_1$.

*We only consider pairs of replica indices (ab) such that $a \neq b$, and the reason for doing so will become clear later in the text.

For B we find

$$B_{(ab)(cd)} = -\alpha^2\beta^3 \left[\left\langle \phi^a\phi^b\phi^c\phi^d \right\rangle_\phi - \left\langle \phi^a\phi^b \right\rangle_\phi \left\langle \phi^c\phi^d \right\rangle_\phi \right] \tag{3.42}$$

$$\langle \cdots \rangle_\phi := \frac{\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}} \cdots e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n} R_{ab}\phi^a\phi^b}}{\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}} e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n} R_{ab}\phi^a\phi^b}} \tag{3.43}$$

In the RS state, B has the same three types of matrix elements as A. To evaluate them we notice that when $\mathsf{R} = \mathsf{R}_{\mathrm{RS}}$

$$\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}} e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n} R_{ab}\phi^a\phi^b} = E_z\left[\left(\int_{-1}^{1} e^{-\beta V_z(\phi)d\phi}\right)^n\right] \simeq E_z\left[1 + n\ \log\int_{-1}^{1} e^{-\beta V_z(\phi)}d\phi\right] \xrightarrow{n\to 0} 1$$

and

$$\begin{aligned}\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}} (\phi^a)^\theta \left(\phi^b\right)^\theta e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n} R_{ab}\phi^a\phi^b} &= \\ = E_z\left[\left(\int_{-1}^{1} \phi^\theta e^{-\beta V_z(\phi)}d\phi\right)^2 \left(\int_{-1}^{1} e^{-\beta V_z(\phi)}d\phi\right)^{n-2}\right] &= E_z\left[\left\langle \phi^\theta \right\rangle_V^2 \left(\int_{-1}^{1} e^{-\beta V_z(\phi)}d\phi\right)^n\right] = \\ \simeq E_z\left\{\left\langle \phi^\theta \right\rangle_V^2 \left[1 + n\ \log\int_{-1}^{1} e^{-\beta V_z(\phi)}d\phi\right]\right\} &\xrightarrow{n\to 0} E_z\left[\left\langle \phi^\theta \right\rangle_V^2\right]\end{aligned}$$

where V_z is given by (3.25) and $\langle\ \rangle_V$ is defined in (3.29). Using these results one easily shows that

$$B_{(ab)(ab)} \equiv B_1 = -\alpha^2\beta^3 \left\{ E_z\left(\langle\phi^2\rangle_V^2\right) - \left[E_z\left(\langle\phi\rangle_V^2\right)\right]^2 \right\} \tag{3.44}$$

$$B_{(ab)(ac)} \equiv B_2 = -\alpha^2\beta^3 \left\{ E_z\left(\langle\phi^2\rangle_V \langle\phi\rangle_V^2\right) - \left[E_z\left(\langle\phi\rangle_V^2\right)\right]^2 \right\} \tag{3.45}$$

$$B_{(ab)(cd)} \equiv B_3 = -\alpha^2\beta^3 \left\{ E_z\left(\langle\phi\rangle_V^4\right) - \left[E_z\left(\langle\phi\rangle_V^2\right)\right]^2 \right\} \tag{3.46}$$

Finally, it is simple to understand that for C one has

$$C_{(ab)(cd)} = \alpha\beta\left(\delta_{ac}\delta_{bd} + \delta_{ad}\delta_{bc}\right) \tag{3.47}$$

We want study the effect of fluctuations around the RS solution:

$$Q_{ab} \to Q_{ab} + \delta Q_{ab} \qquad R_{ab} \to R_{ab} + \delta R_{ab} \tag{3.48}$$

The eigenvalue equations for S is

$$\mathsf{S}\begin{pmatrix}\boldsymbol{\delta Q}\\ \boldsymbol{\delta R}\end{pmatrix} = \lambda\begin{pmatrix}\boldsymbol{\delta Q}\\ \boldsymbol{\delta R}\end{pmatrix} \qquad \boldsymbol{\delta Q} = \{\delta Q_{ab}\}_{a<b}^{1,n} \qquad \boldsymbol{\delta R} = \{\delta R_{ab}\}_{a<b}^{1,n} \tag{3.49}$$

Following the usual practice we take $\delta R_{ab} = x\ \delta Q_{ab}$ and assume that the destabilization is symmetric and affects only the off-diagonal elements of the overlap matrix Q: $\delta Q_{aa} = 0$ and $\delta Q_{ab} = \delta Q_{ba}$. S (like A, B, and C) falls in the class of matrices called *ultrametric* [43]. There are three possible classes of eigenvectors (and correspondingly of eigenvalues), namely those invariant under interchange of all replica indices, all but one replica index and all but two replica indices, respectively. Fluctuations in the longitudinal direction correspond to the first two of these classes and can be neglected whenever a unique saddle point is obtained, as seems to be the case. It can in fact be shown that their corresponding eigenvalues are always positive. The last class of eigenvectors describes the transverse mode, called the "replicon mode" in the disordered systems literature. From this last case we expect problems. We therefore consider fluctuations of the form

$$\delta Q_{ab} = \begin{cases} \frac{1}{2}(3-n)(2-n)x & (a=1, b=2) \\ \frac{1}{2}(3-n)x & (a=1,2) \text{ and } (b\neq 1,2) \\ x & (a,b\neq 1,2) \end{cases} \tag{3.50}$$

which ensures orthogonality between longitudinal and transverse eigenvectors. A short calculation shows that, when $n \to 0$, the corresponding eigenvalue equations reduce to

$$\begin{cases} \lambda_{\mathsf{A}} + \alpha\beta x = \lambda \\ x\lambda_{\mathsf{B}} + \alpha\beta = x\lambda \end{cases} \tag{3.51}$$

where $\lambda_{\mathsf{A}} = A_1 - 2A_2 + A_3 = -\beta/[\alpha(1+\chi)^2]$ and $\lambda_{\mathsf{B}} = B_1 - 2B_2 + B_3 = -\alpha^2\beta^3 E_z[(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2)^2]$. Redefining $\frac{\lambda}{\alpha\beta} \to \lambda$ we obtain

$$\begin{cases} x = \lambda + u \\ \frac{1}{x} = \lambda + v \end{cases} \tag{3.52}$$

where $u = [\alpha^2(1+\chi)^2]^{-1}$ and $v = \alpha\beta^2 E_z[(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2)^2]$. Notice that $u, v > 0$. From this we obtain the eigenvalues

$$\lambda_\pm = -\frac{1}{2}(u+v) \pm \frac{1}{2}\sqrt{(u-v)^2 + 4} \tag{3.53}$$

λ_- never changes sign. In particular, it is always negative, meaning that its sign has to be corrected by a proper choice of the integration contour for R (remember that in obtaining the replica free energy we re-defined R_{ab} as complex numbers). Therefore, the sign of λ_+ controls that of $\det\mathsf{S}$. It is easy to see that $\lambda_+ > 0$ (hence the RS solution is stable) as long as $uv < 1$. The instability sets on when $uv = 1$, i.e. when

$$\lim_{\beta\to\infty} \beta^2 \; E_z\left[\left(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2\right)^2\right] = \alpha(1+\chi)^2 \tag{3.54}$$

To calculate explicitly the average appearing above, we consider the function

$$F_z(h) = \frac{1}{\beta}\log\int_{-1}^{1} e^{\beta W_z(\phi,h)} d\phi \tag{3.55}$$

with $W_z(\phi,h) = -V_z(\phi) + h\phi$, which possesses the remarkable property that $\partial_h^2 F_z(h)|_{h=0} = \beta(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2)$, allowing us to re-write the AT-line as $\lim_{\beta\to\infty} E_z[(\partial_h^2 F_z(h)|_{h=0})^2] = \alpha(1+\chi)^2$. Noting that

$$\begin{aligned} \lim_{\beta\to\infty} F_z(h) &= \lim_{\beta\to\infty}\frac{1}{\beta}\log\left[e^{\beta W_z(\phi^*,h)}\int_{-1}^{1} e^{\beta[W_z(\phi,h)-W_z(\phi^*,h)]}d\phi\right] = \\ &= W_z(\phi^*,h) + \lim_{\beta\to\infty}\frac{1}{\beta}\log\int_{-1}^{1} e^{\beta[W_z(\phi,h)-W_z(\phi^*,h)]}d\phi = W_z(\phi^*,h) \end{aligned} \tag{3.56}$$

where ϕ^* is the maximum of $W_z(\phi,h)$ in $[-1,1]$, we conclude that the AT-line is

$$E_z\left[\left(\partial_h^2 W_z(\phi^*,h)|_{h=0}\right)^2\right] = \alpha(1+\chi)^2 \tag{3.57}$$

Setting $W_z(\phi,h) = -(1/2)b\phi^2 + az\phi + h\phi$, where $b = -\alpha\beta(R-r)$ and $a = \sqrt{\alpha r}$ (the coefficients a and b thus defined are both positive), one easily sees that if $|(az+h)/b| < 1$ then $\phi^* = (az+h)/b$, so $W_z(\phi^*,h) = \frac{1}{2b}(az+h)^2$ and $\partial_h^2 W_z(\phi^*,h)|_{h=0} = 1/b$. On the other hand, if $|(az+h)/b| > 1$, then the root of the derivative of $W_z(\phi,h)$ lies outside the admitted interval, so that $\phi^* = 1$ for $(az+h)/b > 1$ and $\phi^* = -1$ for $(az+h)/b < -1$. In these cases one has, respectively, $W_z(\phi^*,h) = b/2 \pm (az+h)$ so that $\partial_h^2 W_z(\phi^*,h)|_{h=0} = 0$. The Gaussian average in (3.57) therefore splits in two parts,

$$E_z\left[\left(\partial_h^2 W_z(\phi^*,h)|_{h=0}\right)^2\right] = \frac{1}{\sqrt{2\pi}}\int_{|z|<b/a}\frac{1}{b^2}\;e^{-\frac{1}{2}z^2}dz + \frac{1}{\sqrt{2\pi}}\int_{|z|>b/a} 0\times e^{-\frac{1}{2}z^2}dz = \frac{1}{b^2}\mathrm{erf}\left(\frac{b}{a\sqrt{2}}\right)$$

where the second integral above clearly vanishes. Using (3.27) to have b in terms of χ and η, the AT-line reads $\alpha[1-\eta(1+\chi)]^2 = \mathrm{erf}[b/(a\sqrt{2})]$. Dynamical analysis (next chapter) will provide us with a clear interpretation of the right-hand side. For the moment, let it suffice to say that one can show that the AT-line has the final expression

$$[1-\eta(1+\chi)]^2 = (1-\sqrt{\eta})^2 \tag{3.58}$$

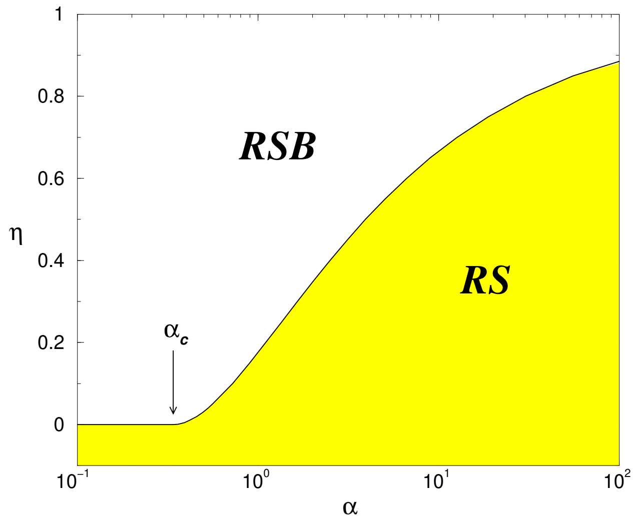

FIGURE 3.3. The AT-line.

This shows, in particular, that replica-symmetry can be broken for all $\eta > 0$. For $\eta \geq 1$, then, RSB occurs at all values of α. The AT-line is reported in Fig. 3.3. Eq. (3.58) describes a set of pairs (α, η) where second-order phase transitions (χ diverges) to a spin-glass state, where ergodicity is broken, takes place.

3.5. Ground-state entropy for $\eta = 1$

One of the most intriguing features to which RSB is connected in the mean-field theory of spin-glasses is the existence of a large number of metastable states. For instance, in the SK model it has been shown that the number of such states is exponentially large in the system size. We want to address an analogous question in the our context. We will concentrate on the $\eta = 1$ model, where RSB occurs at all values of α, since there the calculation simplifies considerably because the 'soft' spins $\phi_i \in [-1, 1]$ can be treated as Ising spins $s_i \in \{-1, 1\}$. As will soon become clear, the $\eta = 1$ case is particularly important from the game-theoretic viewpoint because the metastable states of H_1 are the Nash equilibria of the game.

We start by noting that, upon defining $s_{-i} = \{s_j\}_{j \neq i}^{1,N}$ and

$$u_i^\mu(s_i, s_{-i}) = -a^\mu_{i\widetilde{g}_i} A^\mu = -(\omega_i^\mu + s_i \xi_i^\mu) \sum_{1 \leq j \leq N} (\omega_j^\mu + s_j \xi_j^\mu) \tag{3.59}$$

one has that $H_1(\boldsymbol{s}) \equiv \overline{\langle A^2 \rangle} = -\sum_{1 \leq i \leq N} \overline{u_i^\mu(s_i, s_{-i})}$, where again $\overline{x^\mu} = \sum_{\mu \leq \alpha N} x^\mu / (\alpha N)$. $u_i^\mu(s_i, s_{-i})$ plays the role of an effective local field. One can define a single-agent effective "energy" $s_i \overline{u_i^\mu(s_i, s_{-i})}$ so that the states $\boldsymbol{s}$ which are stable with respect to a single spin flip for H_1 must satisfy the condition

$$s_i \left[\overline{u_i^\mu(+1, s_{-i})} - \overline{u_i^\mu(-1, s_{-i})} \right] \geq 0 \qquad \text{for all } i \tag{3.60}$$

The quantity (3.59) can be interpreted as the payoff to an agent i who is employing strategy $\widetilde{g}_i$ in the strategic context described by $\boldsymbol{s} = \{s_i, s_{-i}\}$. Therefore (3.60) states that $\boldsymbol{s}$ is a strategic configuration such that no agent can improve his expected utility, the expectation being over all possible information patterns, by changing strategy unilaterally if other agents stick to their respective choices. This is the usual definition of a Nash equilibrium in game theory. Counting the number of Nash equilibria of the minority game is thus the same as counting the number of states which are stable with respect to a single spin flip for H_1.

After some simple algebra, (3.60) is found to be equivalent to the request that $1 - 2s_i\overline{\xi_i^\mu A^\mu} \geq 0$ for all i, so that an indicator function is

$$I(\boldsymbol{s}) = \prod_{1\leq i\leq N} \Theta\left(1 - 2s_i\overline{\xi_i^\mu A^\mu}\right) = \begin{cases} 1 & \text{if } \boldsymbol{s} = \{s_i\}_{i=1}^N \text{ is a Nash eq.} \\ 0 & \text{otherwise} \end{cases} \tag{3.61}$$

Θ being the step function, and the number of minima is simply $\mathcal{N} = \sum_{\boldsymbol{s}\in\{-1,1\}^N} I(\boldsymbol{s}) \equiv \mathrm{Tr}_{\boldsymbol{s}} I(\boldsymbol{s})$. Using the integral representation of the Θ function we have

$$\mathcal{N} = \mathrm{Tr}_{\boldsymbol{s}} \int e^{i\sum_{1\leq i\leq N}\widehat{x}_i\left(x_i - 1 + 2s_i\overline{\xi_i^\mu A^\mu}\right)}\, d\boldsymbol{x}\, d\widehat{\boldsymbol{x}} \tag{3.62}$$

$\mathcal{N}$ still depends on the concrete realization $\{\boldsymbol{a}_{ig}\}$ of the quenched disorder. Being interested in typical properties (i.e. valid for almost all realizations of disorder), we should perform some type of disorder average. The correct thing to do at this point would be to calculate the typical number of Nash equilibria from the so-called "quenched approximation", $\mathcal{N}_{\text{typ}} = e^{[\log\mathcal{N}]_{\boldsymbol{a}}}$, because $\mathcal{N}$ is a product of many random contributions, and products of random numbers have distributions with long tails for which the average and the most probable value can be utterly different. The logarithm of $\mathcal{N}$, instead, is a sum of N independent random terms which, for sufficiently large N, becomes Gaussian distributed so that its average and most probable value asymptotically ($N \to \infty$) coincide. We choose however to perform directly the disorder average of $\mathcal{N}$, defining $\mathcal{N}_{\text{typ}} = [\mathcal{N}]_{\boldsymbol{a}}$. This choice corresponds to the so-called "annealed approximation" which, at odds with the quenched calculation, does not require the use of the replica formalism and is considerably more straightforward. We will see that in spite of this flaw the results obtained in this way are already very satisfactory.

The disorder average takes a particularly simple form if we isolate the relevant terms inserting the αN variables $b^\mu = (1/\sqrt{N})A^\mu \quad (\mu = 1,\ldots,\alpha N)$ using the identity

$$1 = \int d\boldsymbol{b} \prod_{\mu\leq\alpha N} \delta\left(b^\mu - \frac{1}{\sqrt{N}}A^\mu\right) = (2\pi)^{-\alpha N}\int e^{i\sum_{\mu\leq\alpha N}\widehat{b}^\mu\left(b^\mu - \frac{1}{\sqrt{N}}A^\mu\right)} d\boldsymbol{b}\, d\widehat{\boldsymbol{b}} \tag{3.63}$$

Recalling that $A^\mu = \frac{1}{2}\sum_{1\leq i\leq N}\left[(a_{i1}^\mu + a_{i2}^\mu) + s_i\,(a_{i1}^\mu - a_{i2}^\mu)\right]$ one easily arrives at

$$\mathcal{N} = \mathrm{Tr}_{\boldsymbol{s}}\int d\boldsymbol{x}\, d\widehat{\boldsymbol{x}}\; e^{i\sum_{1\leq i\leq N}\widehat{x}_i(x_i-1)}\int d\boldsymbol{b}\, d\widehat{\boldsymbol{b}}\; e^{i\sum_{\mu\leq\alpha N}\widehat{b}^\mu b^\mu}\prod_{\substack{1\leq i\leq N\\ \mu\leq\alpha N}} e^{-\frac{i\widehat{b}^\mu}{2\sqrt{N}}\left[a_{i1}^\mu(1+s_i)+a_{i2}^\mu(1-s_i)\right]+\frac{i\widehat{x}_i s_i b^\mu}{\alpha\sqrt{N}}\left(a_{i1}^\mu - a_{i2}^\mu\right)}$$

The disorder average $[\mathcal{N}]_{\boldsymbol{a}}$ now factorizes nicely:

$$[\mathcal{N}]_{\boldsymbol{a}} = \mathrm{Tr}_{\boldsymbol{s}}\int d\boldsymbol{x}\, d\widehat{\boldsymbol{x}}\; e^{i\sum_{1\leq i\leq N}\widehat{x}_i(x_i-1)}\int d\boldsymbol{b}\, d\widehat{\boldsymbol{b}}\; e^{i\sum_{\mu\leq\alpha N}\widehat{b}^\mu b^\mu}\prod_{\substack{1\leq i\leq N\\ \mu\leq\alpha N}}[X_i^\mu]_{\boldsymbol{a}} \tag{3.64}$$

$$X_i^\mu = e^{i\left[\frac{\widehat{x}_i s_i b^\mu}{\alpha\sqrt{N}} - \frac{\widehat{b}^\mu(1+s_i)}{2\sqrt{N}}\right]a_{i1}^\mu}\, e^{-i\left[\frac{\widehat{x}_i s_i b^\mu}{\alpha\sqrt{N}} + \frac{\widehat{b}^\mu(1-s_i)}{2\sqrt{N}}\right]a_{i2}^\mu} \tag{3.65}$$

so that, using the disorder probability density $P(a_{ig}^\mu) = \frac{1}{2}(\delta_{a_{ig}^\mu,+1} + \delta_{a_{ig}^\mu,-1}) \quad (g = 1,2)$ and a little algebra, one obtains

$$[X_i^\mu]_{\boldsymbol{a}} = e^{-\frac{1}{2N}(\widehat{b}^\mu)^2 - \frac{(b^\mu)^2\widehat{x}_i^2}{\alpha^2 N} + \frac{b^\mu\widehat{b}^\mu\widehat{x}_i}{\alpha N}} \tag{3.66}$$

Introducing the order parameters

$$\Gamma = -\frac{1}{\alpha^2 N}\sum_{1\leq i\leq N}\widehat{x}_i^2 \qquad \text{and} \qquad \gamma = \frac{1}{i\alpha N}\sum_{1\leq i\leq N}\widehat{x}_i \tag{3.67}$$

through appropriate δ-distributions with the identity (all integrals are from $-\infty$ to $+\infty$)

$$\int i\alpha N d\gamma\int\alpha^2 N d\Gamma\int\frac{d\omega}{2\pi}\int\frac{d\Omega}{2\pi}\; e^{i\omega\left(i\alpha N\gamma - \sum_{1\leq i\leq N}\widehat{x}_i\right)}\; e^{-i\Omega\left(\alpha^2 N\Gamma + \sum_{1\leq i\leq N}\widehat{x}_i^2\right)} = 1 \tag{3.68}$$

one has

$$[\mathcal{N}]_{\boldsymbol{a}} = \frac{i\alpha^3 N^2}{(2\pi)^2}\mathrm{Tr}_{\boldsymbol{s}} \int d\gamma\ d\Gamma\ d\omega\ d\Omega\ e^{-\alpha N\gamma\omega - i\Omega\alpha^2 N\Gamma} \times \\ \times \int d\boldsymbol{x}\ d\widehat{\boldsymbol{x}}\ e^{i\sum_{1\le i\le N}\left[\widehat{x}_i(x_i-1-\omega)-\Omega\widehat{x}_i^2\right]} \int d\boldsymbol{b}\ d\widehat{\boldsymbol{b}}\ e^{\sum_{\mu\le\alpha N}\left[i(1+\gamma)b^\mu\widehat{b}^\mu+\Gamma(b^\mu)^2-\frac{1}{2}\left(\widehat{b}^\mu\right)^2\right]} \tag{3.69}$$

We can now integrate over $(\boldsymbol{x},\widehat{\boldsymbol{x}})$ and $(\boldsymbol{b},\widehat{\boldsymbol{b}})$, obtaining

$$\int_{-\infty}^{\infty} db \int_{-\infty}^{\infty} \frac{d\widehat{b}}{2\pi}\ e^{\left[i(1+\gamma)b\widehat{b}+\Gamma b^2-\frac{1}{2}\widehat{b}^2\right]} = e^{-\frac{1}{2}\log\left[(1+\gamma)^2-2\Gamma\right]} \tag{3.70}$$

$$\int_{0}^{\infty} dx \int_{-\infty}^{\infty} \frac{d\widehat{x}}{2\pi}\ e^{i\left[\widehat{x}(x-1-\omega)-\Omega\widehat{x}^2\right]} = \frac{1}{2}\left[1+\mathrm{erf}\left(\frac{1+\omega}{2\sqrt{i\Omega}}\right)\right] \tag{3.71}$$

Notice that there are αN integrals of the first type and N of the second type. From the latter we learn that $i\Omega$ is real and positive. Redefining $i\Omega \to \Omega$ and noticing that, as its argument is independent of $\boldsymbol{s}$, $\mathrm{Tr}_{\boldsymbol{s}}$ can be replaced simply by 2^N, we end up with

$$[\mathcal{N}]_{\boldsymbol{a}} = \frac{\alpha^3 N^2}{(2\pi)^2} \int e^{N\Sigma(\gamma,\Gamma,\omega,\Omega)}\ d\gamma\ d\Gamma\ d\omega\ d\Omega \tag{3.72}$$

where

$$\Sigma(\gamma,\Gamma,\omega,\Omega) = -\alpha\gamma\omega - \alpha^2\Gamma\Omega - \frac{\alpha}{2}\log\left[(1+\gamma)^2-2\Gamma\right] + \log\left[1+\mathrm{erf}\left(\frac{1+\omega}{2\sqrt{\Omega}}\right)\right] \tag{3.73}$$

$[\mathcal{N}]_{\boldsymbol{a}}$ is asymptotically dominated by the saddle point of Σ, which is easily seen to be attained at $\omega = 1-\gamma^*$, $\Omega = (1-\gamma^*)/[\alpha(1+\gamma^*)]$, and $\Gamma = -(\gamma^*)^2(1+\gamma^*)/[2(1-\gamma^*)]$, where γ^* is the root of the equation

$$x^2 = \log(2x) - \log\left\{\alpha\gamma^2\sqrt{\pi}\left[1+\mathrm{erf}(x)\right]\right\} \qquad x = \frac{\gamma}{2}\sqrt{\frac{\alpha(1+\gamma)}{1-\gamma}} \tag{3.74}$$

In terms of γ^* we have

$$\Sigma(\alpha) = \frac{\alpha\gamma^*}{2}(2-\gamma^*) - \frac{\alpha}{2}\log\left(\frac{1+\gamma^*}{1-\gamma^*}\right) + \log\left[1+\mathrm{erf}\left(\frac{\gamma^*}{2}\sqrt{\frac{\alpha(1+\gamma^*)}{1-\gamma^*}}\right)\right] \tag{3.75}$$

which is our final expression for the annealed ground-state entropy of the $\eta = 1$ model. In Fig. 3.4. one sees the analytic prediction compared to numerical results obtained via exact enumeration of metastable states. The agreement is satisfactory.

3.6. Replica-symmetry breaking

We will now study the effects of RSB on the solution of the minimization problem for H_η. This will provide us with a better description of the non-ergodic regime (where replica-symmetric predictions deviate from computational results). The general expression for the free energy, Eq. (3.21), has to be evaluated under the hypothesis that the overlap matrix Q has a more complicated structure that the simple replica-symmetric one, in particular one which breaks the permutation symmetry between replicas. The simplest possibility is given by Parisi's one-step RSB Ansatz (1RSB), expressed by (2.16):

$$\mathsf{Q} \equiv \mathsf{Q}_{\mathrm{1RSB}} = q_0\ \mathsf{E}_n + (q_1-q_0)\ \mathsf{I}_{\frac{n}{m}} \otimes \mathsf{E}_m + (Q-q_1)\ \mathsf{I}_n, \tag{3.76}$$

where we have explicitly written the dimensionality of the involved matrices as subscripts, and $\otimes$ denotes the standard tensor (Kronecker) product. In other terms,

$$Q_{ab} = \begin{cases} Q & \text{for } a = b \\ q_1 & \text{for } a \neq b \text{ such that } I\left(\frac{a}{m}\right) = I\left(\frac{b}{m}\right) \\ q_0 & \text{otherwise} \end{cases} \tag{3.77}$$

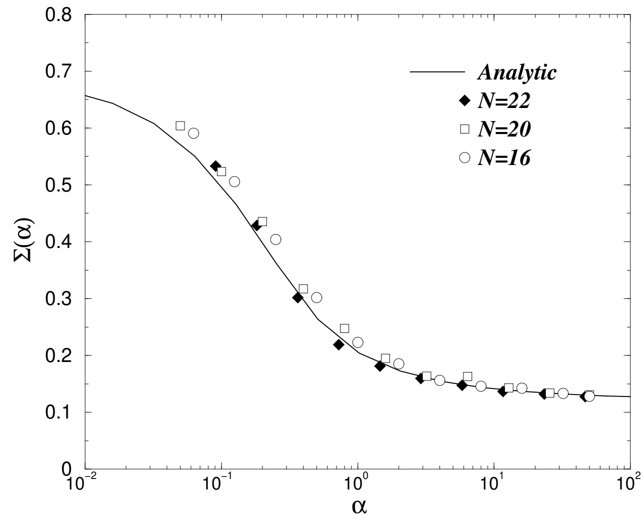

FIGURE 3.4. Ground-state entropy of the $\eta = 1$ model. Analytic results obtained via annealed approximation. Numerical results correspond to exact enumeration of Nash equilibria, defined by condition (3.60) for systems of size $N = 16, 20, 22$.

Here m is a partitioning parameter that is to be determined self-consistently from a stationarity condition for the free energy, and $I(x)$ is the integer-valued function, equal to the smallest integer not smaller than x. Clearly, an analogous Ansatz has to be made for the matrix R of Lagrange multipliers.

We shall now evaluate all terms of the free energy. As before, we need to consider the matrix T, now given by

$$\mathsf{T} = \left[1 + \frac{\beta}{\alpha}(Q - q_1)\right] \; \mathsf{I}_n + \frac{\beta}{\alpha}(1 + q_0) \; \mathsf{E}_n + \frac{\beta}{\alpha}(q_1 - q_0) \; \mathsf{I}_{\frac{n}{m}} \otimes \mathsf{E}_m \tag{3.78}$$

for which we get the following spectrum ($k = 1, \dots, n/m$):

Eigenspace :	Eigenvalue :	Degeneracy :
$\sum_{a=1}^{m} x_{(k-1)m+a} = 0 \;,\; \forall k$	$\lambda_1 = 1 + \beta(Q - q_1)/\alpha$	$\frac{n}{m}(m-1) = n - \frac{n}{m}$
$\sum_{k=1}^{n/m}\sum_{a=1}^{m} x_{(k-1)m+a} = 0$	$\lambda_2 = \lambda_1 + m\beta(q_1 - q_0)/\alpha$	$\frac{n}{m} - 1$
$\boldsymbol{x} = (1, \dots, 1)$	$\lambda_3 = \lambda_2 + n\beta(1 + q_0)/\alpha$	1

so that[†]

$$\lim_{n\to 0}\frac{1}{n}\log\det\mathsf{T} = \lim_{n\to 0}\frac{1}{n}\log\left[\lambda_1^{n-\frac{n}{m}}\lambda_2^{\frac{n}{m}-1}\lambda_3\right] = \tag{3.79}$$
$$= \log\left[1 + \frac{\beta}{\alpha}(Q - q_1)\right] + \frac{\beta(1+q_0)}{\alpha + \beta(Q-q_1) + m\beta(q_1 - q_0)} + \frac{1}{m}\log\left[1 + \frac{m\beta(q_1 - q_0)}{\alpha + \beta(Q - q_1)}\right]$$

For the remaining terms one has

$$\lim_{n\to 0}\frac{1}{n}\sum_{a,b=1}^{n} R_{ab}Q_{ab} = \lim_{n\to 0}\frac{1}{n}\left[nRQ + \frac{n}{m}m(m-1)r_1q_1 + \left(n^2 - \frac{n}{m}m^2\right)r_0q_0\right] = \tag{3.80}$$
$$= RQ + (m-1)r_1q_1 - mr_0q_0$$

[†]One obtains the same result in a more "direct" way using the identities $\mathsf{I}_n \equiv \mathsf{I}_{\frac{n}{m}} \otimes \mathsf{I}_m$ and $\mathsf{E}_n \equiv \mathsf{E}_{\frac{n}{m}} \otimes \mathsf{E}_m$ and considering that matrices with tensor product structure can be diagonalized separately in each tensor product component.

and

$$\lim_{n\to 0}\frac{1}{n}\log\left[\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}}e^{\frac{\alpha\beta^2}{2}\sum_{1\le a,b\le n}R_{ab}\phi^a\phi^b}\right]=$$

$$=\lim_{n\to 0}\frac{1}{n}\log\left\{\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}}e^{\frac{\alpha\beta^2}{2}\left[(R-r_1)\sum_{1\le a\le n}(\phi^a)^2+(r_1-r_0)\sum_{k\le n/m}\left(\sum_{a\le m}\phi^{(k-1)m+a}\right)^2+r_0\left(\sum_{1\le a\le n}\phi^a\right)^2\right]}\right\}=$$

$$=\lim_{n\to 0}\frac{1}{n}\log\ E_z\left\{\prod_{k\le n/m}E_{y_k}\left[\mathrm{Tr}_{\{\phi^1,\dots,\phi^n\}}e^{z\beta\sqrt{\alpha r_0}\sum_{1\le a\le n}\phi^a+\beta\sqrt{\alpha(r_1-r_0)}\sum_{k\le n/m}y_k\left(\sum_{a\le m}\phi^{(k-1)m+a}\right)+\frac{\alpha\beta^2}{2}(R-r_1)\sum_{1\le a\le n}(\phi^a)^2}\right]\right\}=$$

$$=\lim_{n\to 0}\frac{1}{n}\log\ E_z\left[E_y\left(\int_{-1}^{1}e^{-\beta V_{yz}(\phi)}d\phi\right)^m\right]^{\frac{n}{m}}\simeq\lim_{n\to 0}\frac{1}{n}\log\ E_z\left[1+\frac{n}{m}\log E_y\left(\int_{-1}^{1}e^{-\beta V_{yz}(\phi)}d\phi\right)^m\right]=$$

$$=\frac{1}{m}\ E_z\left\{\log E_y\left[\left(\int_{-1}^{1}e^{-\beta V_{yz}(\phi)}d\phi\right)^m\right]\right\}$$

where

$$V_{yz}(\phi)=-\sqrt{\alpha r_0}z\phi-\sqrt{\alpha(r_1-r_0)}y\phi-\frac{\alpha\beta}{2}(R-r_1)\phi^2 \tag{3.81}$$

Putting fragments together one has

$$\begin{aligned}f(\mathsf{Q}_{1\mathrm{RSB}},\mathsf{R}_{1\mathrm{RSB}})&=\frac{\alpha}{2\beta}\log\left[1+\frac{\beta}{\alpha}(Q-q_1)\right]+\frac{\alpha}{2\beta m}\log\left[1+\frac{m\beta(q_1-q_0)}{\alpha+\beta(Q-q_1)}\right]+\\&+\frac{\alpha}{2}\frac{1+q_0}{\alpha+\beta(Q-q_1)+m\beta(q_1-q_0)}+\frac{\alpha\beta}{2}[RQ+(m-1)r_1q_1-mr_0q_0]+\frac{\eta}{2}(1-Q)+\\&-\frac{1}{m\beta}\ E_z\left\{\log\ E_y\left[\left(\int_{-1}^{1}e^{-\beta V_{yz}(\phi)}d\phi\right)^m\right]\right\}\end{aligned} \tag{3.82}$$

By variating of f with respect to the overlaps Q, q_0 and q_1, to their conjugate variables, and to m one obtains seven saddle-point equations, which after some algebraic manipulation take the form

$$Q=\langle\!\langle\phi^2\rangle\!\rangle \qquad r_1-r_0=\frac{q_1-q_0}{[\alpha+\beta(Q-q_1)][\alpha+\beta(Q-q_1)+\beta m(q_1-q_0)]}$$

$$\beta m(q_1-q_0)+\beta(Q-q_1)=\frac{1}{\sqrt{\alpha r_0}}\langle\!\langle z\phi\rangle\!\rangle \qquad \beta(R-r_1)=\frac{\eta}{\alpha}-\frac{1}{\alpha+\beta(Q-q_1)}$$

$$r_0=\frac{1+q_0}{[\alpha+\beta(Q-q_1)+\beta m(q_1-q_0)]^2} \qquad \beta m q_1+\beta(Q-q_1)=\frac{1}{\sqrt{\alpha(r_1-r_0)}}\langle\!\langle y\phi\rangle\!\rangle$$

$$\frac{1}{\beta}E_z\left[\frac{E_y[\mathcal{Z}^m\log\mathcal{Z}]}{E_y[\mathcal{Z}^m]}\right]-\frac{1}{\beta m}E_z[\log[E_y(\mathcal{Z}^m)]]=\frac{\alpha}{2}\frac{(q_1-q_0)[\alpha+\beta(Q-q_1)+\beta m q_1]}{[\alpha+\beta(Q-q_1)][\alpha+\beta(Q-q_1)+\beta m(q_1-q_0)]}+$$
$$-\frac{\alpha}{2\beta m}\log\left[1+\frac{\beta m(q_1-q_0)}{\alpha+\beta(Q-q_1)}\right]$$

Here, we have set $\mathcal{Z}=\int_{-1}^{1}e^{-\beta V_{yz}(\phi)}d\phi$ and

$$\langle\!\langle\cdots\rangle\!\rangle=E_z\left[\frac{E_y[\mathcal{Z}^{m-1}\int_{-1}^{1}\cdots e^{-\beta V_{yz}(\phi)}d\phi]}{E_y[\mathcal{Z}^m]}\right] \tag{3.83}$$

The analysis of these equations is quite involved even in the limit $\beta\to\infty$, and we won't report it here. Let it suffice to mention that numerical investigations show three different regimes in the (α,η) plane when $\beta\to\infty$ (see Fig.):

- For $\alpha<\alpha_0\simeq 0.09012\ldots$ (all $\eta>0$) one has $H_\eta=0$. The solution does not depend on η as long as $\eta>0$. The self-overlap is $Q=1$ signalling that agents play pure strategies ($\phi_i=\pm 1$) but off diagonal overlaps $q_1>q_0$ are both less than 1. This suggests that NE are organized in a complex geometric structure. The parameter m attains a finite value.
- For $\alpha_0<\alpha<\alpha_1(\eta)$ (all $\eta>0$) the solution has $H_\eta>0$, it is independent of η (for $\eta>0$) and $1=Q=q_1>q_0$. The spin susceptibility $\chi=\beta(Q-q_1)/\alpha$ attains a finite value in the limit $\beta\to\infty$, which diverges as $\alpha\to\alpha_0$. Again agents play pure strategies and $q_0<1$ is the

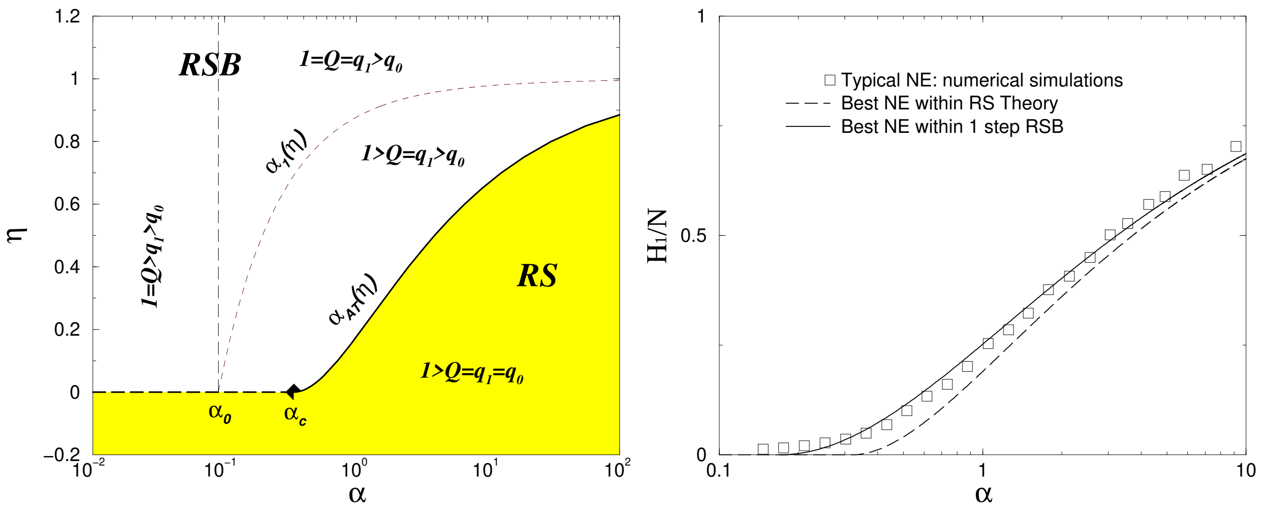

(a) (b)

FIGURE 3.5. (a) Phase diagram of the minority game within the 1RSB Ansatz. (b) Comparison between the typical value of H_1/N (markers), the replica-symmetric approximation (dashed line), and the one-step replica-symmetry breaking result (continuous line).

typical overlap between two NE. The parameter m vanishes as $1/\beta$ (indeed the βm is finite as $\beta \to \infty$). The line $\alpha_1(\eta)$ is determined by the solution of

$$\frac{\eta}{2} = \frac{1}{\alpha + \beta(Q - q_1)}. \tag{3.84}$$

- In between the line $\alpha_1(\eta)$ and the stability line $\alpha_{\rm AT}(\eta)$ the solution has $H_\eta > 0$ and $1 > Q = q_1 > q_0$. Hence agents do not play pure strategies. The solution in this region depends on η.

Figure 3.5 shows that the one-step calculation for H_1/N agrees very well with numerical simulations and it represents a considerable improvement over the replica symmetric result‡. Further steps of RSB, most probably infinitely many, are likely to be needed to recover the exact solution. However, already the one step calculation provides a rather good approximation.

‡Note that numerical results refer to a typical NE which need not be the ground state of H_1.

CHAPTER 4

Dynamics

In this chapter we solve the dynamics of the minority game using generating functional techniques, to carry out the disorder-average explicitly. First, we introduce a modification in the model's dynamics, which does not change the qualitative behaviour but simplifies the calculations considerably. Then we will evaluate the generating functional and solve ergodic stationary states upon making customary assumptions. We will show that ergodicity breaking in this model is not related to a breakdown of time-translation invariance but, rather, to a failure of weak long-term memory. This makes a solution of the dynamics in the non-ergodic regime a hopeless task. However, we calculate a condition for the onset of long-term memory, finding that it coincides with the AT-line derived in Chapter 2. We support these results with numerical simulations. Finally, we discuss the physical reasons for this unusual picture.

4.1. 'Batch' minority game

Equilibrium techniques have provided us with much detailed quantitative information about the stationary behaviour of the system in the limit $N \to \infty$. We address now the question of the approach to stationarity by means of dynamical techniques.

We begin with some notation. The total bid $A(n)$ as defined in (5.10) is $\mathcal{O}(\sqrt{N})$. It is convenient to re-scale it to make it $\mathcal{O}(1)$:

$$A(n) = \frac{1}{\sqrt{N}} \sum_{1 \le i \le N} b_i(n) \tag{4.1}$$

In terms of this quantity the dynamics (2.4) reads

$$p_{ig}(n+1) = p_{ig}(n) - a_{ig}^{\mu(n)} \left[A(n) - \frac{\eta}{\sqrt{N}} \left(a_{i\tilde{g}_i(n)}^{\mu(n)} - a_{ig}^{\mu(n)} \right) \right] \tag{4.2}$$

Correspondingly, the time evolution of the preferences y_i is given by

$$\begin{aligned} y_i(n+1) = y_i(n) - \xi_i^{\mu(n)} \left[A(n) - \frac{\eta}{\sqrt{N}}\, a_{i\tilde{g}_i(n)}^{\mu(n)} \right] = \\ = y_i(n) - \xi_i^{\mu(n)} \left\{ \Omega^{\mu(n)} + \frac{1}{\sqrt{N}} \sum_{1 \le j \le N} \xi_j^{\mu(n)} s_j(n) - \frac{\eta}{\sqrt{N}} \left[\omega_i^{\mu(n)} + s_i(n) \xi_i^{\mu(n)} \right] \right\} = \\ = y_i(n) - \xi_i^{\mu(n)} \left[\Omega^{\mu(n)} + \frac{1}{\sqrt{N}} \sum_{1 \le j \le N} \xi_j^{\mu(n)} s_j(n) - \frac{\eta}{\sqrt{N}} s_i(n) \xi_i^{\mu(n)} \right] \end{aligned} \tag{4.3}$$

where now $\boldsymbol{\Omega} = (1/\sqrt{N}) \sum_{1 \le i \le N} \boldsymbol{\omega}_i$ and as before $s_i(n) = \text{sign}[y_i(n)]$.

We choose to study a batch version of this model, which is obtained by averaging (4.3) over information patterns:

$$y_i(n+1) = y_i(n) - \frac{1}{\alpha N} \sum_{\mu \le \alpha N} \xi_i^{\mu} \left[\Omega^{\mu} + \frac{1}{\sqrt{N}} \sum_{1 \le j \le N} \xi_j^{\mu} s_j(n) - \frac{\eta}{\sqrt{N}} s_i(n) \xi_i^{\mu} \right] \tag{4.4}$$

giving

$$y_i(t+1) = y_i(t) - \sum_{1 \le j \le N} J_{ij} s_j(t) - h_i + \eta \alpha s_i(t) \tag{4.5}$$

where the 'external field' and 'couplings' are given by $h_i = (2/\sqrt{N})\boldsymbol{\Omega}\cdot\boldsymbol{\xi}_i$ and $J_{ij} = (2/N)\boldsymbol{\xi}_i\cdot\boldsymbol{\xi}_j$, respectively. Notice that the time unit has been changed accordingly from n to t, which is proportional to the number of updates. The specific choice of the time re-scaling has been made for later convenience.

We borrowed the term 'batch' from the statistical mechanical theory of learning in perceptrons, where one considers the following problem. A 'student' perceptron operates the rule $S:\{-1,1\}^N \to \{-1,1\}$ defined by $S(\boldsymbol{\xi}) = \mathrm{sign}(\boldsymbol{J}\cdot\boldsymbol{\xi})$, parametrized by a weight vector $\boldsymbol{J}\in\mathbb{R}^N$, to associate to each input vector $\boldsymbol{\xi}$ a binary output, trying to emulate a 'teacher' perceptron who operates an analogous rule $T:\{-1,1\}^N \to \{-1,1\}$ given by $T(\boldsymbol{\xi}) = \mathrm{sign}(\boldsymbol{B}\cdot\boldsymbol{\xi})$, characterized by a *fixed* weight vector $\boldsymbol{B}\in\mathbb{R}^N$. In order to improve its performance, the student perceptron modifies its weight vector according to an iterative procedure, using examples of input vectors (or 'questions'), drawn at random from a fixed training set $D = \{\boldsymbol{\xi}^\mu\}_{\mu\le\alpha N}$ (where $\boldsymbol{\xi}^\mu\in\{-1,1\}^N$ for all μ, and each component ξ_i^μ is chosen randomly and independently from $\{-1,1\}$ with uniform probability and fixed, for all $1\le i\le N$ and $\mu\le\alpha N$), and the corresponding values of the teacher outputs $T(\boldsymbol{\xi}^\mu)$. The goal of the student is essentially that of maximizing his overlap $(\boldsymbol{J}\cdot\boldsymbol{B})$ with the teacher. One considers two types of learning, called, respectively, *on-line* and *batch*:

$$\begin{aligned}
&\text{on-line:} \qquad \boldsymbol{J}(n+1) = \boldsymbol{J}(n) + \boldsymbol{F}[\boldsymbol{\xi}^{\mu(n)}, \boldsymbol{J}(n), \boldsymbol{B}]\\
&\text{batch} \qquad\ \ \boldsymbol{J}(n+1) = \boldsymbol{J}(n) + \overline{\boldsymbol{F}[\boldsymbol{\xi}^{\mu}, \boldsymbol{J}(n), \boldsymbol{B}]}
\end{aligned}$$

Here $n = 1,2,\dots$ labels the iteration steps, whereas the bar denotes an average over input vectors: $\overline{x^\mu} = (\alpha N)^{-1}\sum_{\mu\le\alpha N} x^\mu$. $\boldsymbol{F}$ is a function which incorporates all details about the learning process. In the first case, at each stage n an input vector $\boldsymbol{\xi}^{\mu(n)}$ is chosen at random with uniform probability from D and subsequently the weight vector $\boldsymbol{J}$ is updated. One sees that this process is stochastic (indeed, Markovian). In batch learning, instead, the modification that would have been made in the on-line version is averaged over the input vectors in the training set D at each iteration step. The resulting process is therefore a deterministic iterative map.

It is now evident that the original dynamics of the minority game, namely (4.3), where the y_i's are updated after every iteration step, is analogous to on-line learning. In contrast, in the 'batch' case (4.5) the updates are made on the basis of the average effect of all possible choices of the information pattern. A connection to the original model is made if one interprets (4.5) as the cumulated effect of $\mathcal{O}(N)$ iterations of (4.3). The batch dynamics is not exactly equivalent to (4.3), not even for $N\to\infty$, but it yields results which are qualitatively very similar (if not identical) to those of the original model. There are good reasons for turning from the on-line model to the batch version. The theoretical advantage is that one can apply the dynamical techniques directly to (4.5), bypassing the (difficult) problem of constructing a proper continuous-time limit. The numerical advantage is that with the batch dynamics one can simulate larger systems for longer times at the same computational cost.

4.2. The generating functional

We start from the 'batch' dynamics (4.5) with an added external field:

$$y_i(t+1) = y_i(t) - \sum_{1\le j\le N} J_{ij}s_j(t) - h_i + \eta\alpha s_i(t) + \theta_i(t) \tag{4.6}$$

One sees that the transition probability density for going to the preference profile $\boldsymbol{y} = \{y_i\}_{i=1}^N$ from $\boldsymbol{y'}$ is simply $W(\boldsymbol{y}|\boldsymbol{y'}) = \prod_{1\le i\le N}\delta(y_i - y_i' + h_i + \sum_{1\le j\le N} J_{ij}s_j' - \eta\alpha s_i')$, that is, using the integral representation of the δ function,

$$W(\boldsymbol{y}|\boldsymbol{y'}) = (2\pi)^{-N}\int e^{\mathrm{i}\sum_{1\le i\le N}\widehat{y}_i\left(y_i - y_i' + h_i + \sum_{1\le j\le N} J_{ij}s_j' - \eta\alpha s_i'\right)} d\widehat{\boldsymbol{y}} \tag{4.7}$$

with $s_i' = \text{sign}[y_i']$. Inserting this into the generating functional

(4.8)
$$Z[\boldsymbol{\psi}] = \left\langle e^{\mathrm{i}\sum_t \sum_{1\le i\le N}\psi_i(t)y_i(t)}\right\rangle_{\text{paths}} = \int e^{\mathrm{i}\sum_t \sum_{1\le i\le N}\psi_i(t)y_i(t)}\ p[\boldsymbol{y}(0)]\prod_t \left[W(\boldsymbol{y}(t+1)|\boldsymbol{y}(t))\ d\boldsymbol{y}(t)\right]$$

(where $\boldsymbol{y_0} = \boldsymbol{y}(0)$ and $p[\boldsymbol{y}(0)]$ is the probability density of the initial condition) together with the definitions of h_i and J_{ij}, and introducing the quantities

$$x_t^\mu = \sqrt{\frac{2}{N}}\sum_{1\le i\le N}\widehat{s}_i(t)\xi_i^\mu \qquad \text{and} \qquad w_t^\mu = \sqrt{\frac{2}{N}}\sum_{1\le i\le N}\widehat{y}_i(t)\xi_i^\mu \tag{4.9}$$

through appropriate δ functions with conjugate variables $\widehat{x}_t^\mu$ and $\widehat{w}_t^\mu$, respectively, after some algebra one finds

$$\begin{aligned}Z[\boldsymbol{\psi}] = \int D\boldsymbol{w}\ D\widehat{\boldsymbol{w}}\ D\boldsymbol{x}\ D\widehat{\boldsymbol{x}}\ &e^{\mathrm{i}\sum_t\sum_{\mu\le\alpha N}\left[w_t^\mu\widehat{w}_t^\mu + x_t^\mu\widehat{x}_t^\mu + \left(\Omega^\mu/\tau + x_t^\mu\right)w_t^\mu\right]}\times\\ &\times\int p[\boldsymbol{y}(0)]\ e^{-\mathrm{i}\sqrt{\frac{2}{N}}\sum_{1\le i\le N}\sum_{\mu\le\alpha N}\xi_i^\mu\sum_t\left[\widehat{w}_t^\mu\widehat{y}_i(t) + \widehat{x}_t^\mu s_i(t)\right]}\times\\ &\qquad\times e^{\mathrm{i}\sum_t\sum_{1\le i\le N}\{\widehat{y}_i(t)[y_i(t+1)-y_i(t)-\eta\alpha s_i(t)-\theta_i(t)]+\psi_i(t)y_i(t)\}}\ D\boldsymbol{y}\ D\widehat{\boldsymbol{y}}\end{aligned}\tag{4.10}$$

where we have used the shorthands $D\boldsymbol{w} = \prod_t\prod_{\mu\le\alpha N}(dw_t^\mu/\sqrt{2\pi})$, $D\boldsymbol{x} = \prod_t\prod_{\mu\le\alpha N}(dx_t^\mu/\sqrt{2\pi})$ and $D\boldsymbol{y} = \prod_t\prod_{1\le i\le N}[dy_i(t)/\sqrt{2\pi}]$, and similar ones for $\widehat{\boldsymbol{w}}$, $\widehat{\boldsymbol{x}}$ and $\widehat{\boldsymbol{y}}$.

The disorder average involves only the terms where ξ_i^μ and Ω^μ appear. Using the definition of Ω^μ and factorizing when possible one easily finds

$$\begin{aligned}&\left[e^{\mathrm{i}\sqrt{2}\sum_t\sum_{\mu\le\alpha N}w_t^\mu\Omega^\mu - \mathrm{i}\sqrt{\frac{2}{N}}\sum_{1\le i\le N}\sum_{\mu\le\alpha N}\xi_i^\mu\sum_t\left[\widehat{w}_t^\mu\widehat{y}_i(t)+\widehat{x}_t^\mu s_i(t)\right]}\right]_{\boldsymbol{a}} =\\ &\qquad = e^{-\frac{1}{2}\sum_{t,t'}\sum_{\mu\le\alpha N}\left(w_t^\mu w_{t'}^\mu + \widehat{w}_t^\mu L_{tt'}\widehat{w}_{t'}^\mu + 2\widehat{x}_t^\mu K_{tt'}\widehat{w}_{t'}^\mu + \widehat{x}_t^\mu C_{tt'}\widehat{x}_{t'}^\mu\right)}\end{aligned}\tag{4.11}$$

where we have introduced the order parameters

$$C_{tt'} = \frac{1}{N}\sum_{1\le i\le N}s_i(t)s_i(t') \qquad K_{tt'} = \frac{1}{N}\sum_{1\le i\le N}s_i(t)\widehat{y}_i(t') \qquad L_{tt'} = \frac{1}{N}\sum_{1\le i\le N}\widehat{y}_i(t)\widehat{y}_i(t') \tag{4.12}$$

Enforcing the above into the generating functional definition by appropriate δ functions and one gets

$$\begin{aligned}[Z[\boldsymbol{\psi}]]_{\boldsymbol{a}} = \int D\mathsf{C}\ D\widehat{\mathsf{C}}\ D\mathsf{K}\ D\widehat{\mathsf{K}}\ D\mathsf{L}\ D\widehat{\mathsf{L}}\ &\int D\boldsymbol{w}\ D\widehat{\boldsymbol{w}}\ D\boldsymbol{x}\ D\widehat{\boldsymbol{x}}\ e^{\mathrm{i}\sum_t\sum_{\mu\le\alpha N}\left(w_t^\mu\widehat{w}_t^\mu + x_t^\mu\widehat{x}_t^\mu + w_t^\mu x_t^\mu\right)}\times\\ &\times e^{-\frac{1}{2}\sum_{t,t'}\sum_{\mu\le\alpha N}\left(w_t^\mu w_{t'}^\mu + \widehat{w}_t^\mu L_{tt'}\widehat{w}_{t'}^\mu + 2\widehat{x}_t^\mu K_{tt'}\widehat{w}_{t'}^\mu + \widehat{x}_t^\mu C_{tt'}\widehat{x}_{t'}^\mu\right)}\times\\ \times\int p[\boldsymbol{y}(0)]\ &e^{\mathrm{i}\sum_t\sum_{1\le i\le N}\{\widehat{y}_i(t)[y_i(t+1)-y_i(t)-\eta\alpha s_i(t)-\theta_i(t)]+\psi_i(t)y_i(t)\}}\times\\ \times e^{\mathrm{i}\sum_{t,t'}\widehat{L}_{tt'}\left[NL_{tt'}-\sum_{1\le i\le N}\widehat{y}_i(t)\widehat{y}_i(t')\right]+i\sum_{t,t'}\widehat{C}_{tt'}\left[NC_{tt'}-\sum_{1\le i\le N}s_i(t)s_i(t')\right]+i\sum_{t,t'}\widehat{K}_{tt'}\left[NK_{tt'}-\sum_{1\le i\le N}s_i(t)\widehat{y}_i(t')\right]}&\ D\boldsymbol{y}\ D\widehat{\boldsymbol{y}}\end{aligned}$$

where $D\mathsf{C} = \prod_{t,t'}(dC_{tt'}/\sqrt{2\pi})$, $D\widehat{\mathsf{C}} = \prod_{t,t'}(d\widehat{C}_{tt'}/\sqrt{2\pi})$, and analogous shorthands for the other integration variables are meant. Assuming that $p[\boldsymbol{y}(0)] = \prod_{1\le i\le N}p[y_i(0)]$, with some rearrangements one finally arrives at

$$[Z[\boldsymbol{\psi}]]_{\boldsymbol{a}} = \int e^{N(\Psi+\Phi+\Omega)}\ D\mathsf{C}\ D\widehat{\mathsf{C}}\ D\mathsf{K}\ D\widehat{\mathsf{K}}\ D\mathsf{L}\ D\widehat{\mathsf{L}} \tag{4.13}$$

where the three distinct contributions to the exponent have separate and clear meanings. The first one, $\Psi \equiv \Psi(\mathsf{C},\widehat{\mathsf{C}},\mathsf{K},\widehat{\mathsf{K}},\mathsf{L},\widehat{\mathsf{L}})$, is a term that links the two-time order parameters to their conjugates:

$$\Psi = \mathrm{i}\sum_{t,t'}\left(L_{tt'}\widehat{L}_{tt'} + C_{tt'}\widehat{C}_{tt'} + K_{tt'}\widehat{K}_{tt'}\right) \tag{4.14}$$

The second, $\Phi \equiv \Phi(\mathsf{C},\mathsf{K},\mathsf{L})$, contains a description of the statistical properties of the agents' strategies:

$$\Phi = \alpha\ \log \int e^{\mathrm{i}\sum_t(w_t\widehat{w}_t + x_t\widehat{x}_t + w_t x_t) - \frac{1}{2}\sum_{t,t'}(w_t w_{t'} + \widehat{w}_t L_{tt'}\widehat{w}_{t'} + 2\widehat{x}_t K_{tt'}\widehat{w}_{t'} + \widehat{x}_t C_{tt'}\widehat{x}_{t'})}\ Dw\ D\widehat{w}\ Dx\ D\widehat{x} \tag{4.15}$$

The third one, $\Omega \equiv \Omega(\widehat{\mathsf{C}},\widehat{\mathsf{K}},\widehat{\mathsf{L}})$, contains the evolution of the agents' preferences and the generating fields $\boldsymbol{\psi}$, and is the one that will ultimately give us the effective single-agent process:

$$\Omega = \frac{1}{N}\sum_{1\le i\le N}\log\int e^{\mathrm{i}\sum_t\{\widehat{y}(t)[y(t+1)-y(t)-\eta\alpha s(t)-\theta_i(t)]+\psi_i(t)y(t)\}}\ \times$$
$$\times\ p[y(0)]\ e^{-\mathrm{i}\sum_{t,t'}\left[s(t)\widehat{C}_{tt'}s(t')+s(t)\widehat{K}_{tt'}\widehat{y}(t')+\widehat{y}(t)\widehat{L}_{tt'}\widehat{y}(t')\right]}\ Dy\ D\widehat{y} \tag{4.16}$$

According to the theory outlined in Chapter 2, the site-averaged and disorder-averaged auto-correlation and response functions can be obtained asymptotically (i.e. in the limit $N\to\infty$) by taking suitable functional derivatives of (4.13). One may verity that, at the relevant saddle-point,

$$C_{tt'} = \lim_{N\to\infty}\frac{1}{N}\sum_{1\le i\le N}\Big[\big\langle s_i(t)s_i(t')\big\rangle_{\mathrm{paths}}\Big]_{\boldsymbol{a}} \qquad -\mathrm{i}K_{tt'} = \lim_{N\to\infty}\frac{1}{N}\sum_{1\le i\le N}\Big[\big\langle s_i(t)s_i(t')\big\rangle_{\mathrm{paths}}\Big]_{\boldsymbol{a}} \tag{4.17}$$

The second identity above suggests to re-define

$$G_{tt'} = -\mathrm{i}K_{tt'} \tag{4.18}$$

A slightly longer calculation is instead needed to show that, at the same saddle point,

$$L_{tt'} = \lim_{N\to\infty}\frac{1}{N}\sum_{1\le i\le N}\frac{\partial^2\,[Z[\mathbf{0}]]_{\boldsymbol{a}}}{\partial\theta_i(t)\partial\theta_i(t')} = 0 \tag{4.19}$$

4.3. Saddle-point equations and the effective single-agent process

4.3.1. Saddle-point equations. Having served their purpose, all generating fields $\psi_i(t)$ may be put to zero. Moreover, we let $\theta_i(t) = \theta(t)$ for all i. We will now perform the saddle-point integration by taking derivatives of (4.13) with respect to $\{C_{tt'},\widehat{C}_{tt'},K_{tt'},\widehat{K}_{tt'},L_{tt'},\widehat{L}_{tt'}\}$. One finds

$$C_{tt'} = \big\langle s(t)s(t')\big\rangle_* \qquad G_{tt'} = \frac{\partial}{\partial\theta(t')}\langle s(t)\rangle_* \qquad L_{tt'} = 0 \tag{4.20}$$

$$\widehat{C}_{tt'} = \mathrm{i}\frac{\partial\Phi}{\partial C_{tt'}} \qquad \widehat{K}_{tt'} = \mathrm{i}\frac{\partial\Phi}{\partial K_{tt'}} \qquad \widehat{L}_{tt'} = \mathrm{i}\frac{\partial\Phi}{\partial L_{tt'}} \tag{4.21}$$

where $\langle\ \rangle_*$ denotes an effective single-agent average defined as

$$\langle\cdots\rangle_* = \frac{\int\cdots M(y)Dy}{\int M(y)Dy} \tag{4.22}$$

with

$$M(y) = p[y(0)]\ e^{-\mathrm{i}\sum_{t,t'}s(t)\widehat{C}_{tt'}s(t')}\int e^{-\mathrm{i}\sum_{t,t'}\widehat{y}(t)\widehat{L}_{tt'}\widehat{y}(t')+\mathrm{i}\sum_t\widehat{y}(t)\left[y(t+1)-y(t)-\alpha\eta s(t)-\theta(t)-\sum_{t'}\widehat{K}^T_{tt'}s(t')\right]}\ D\widehat{y}$$

Notice that we are now in the position of obtaining exact asymptotic closed equations for the disorder-averaged auto-correlation and response functions in terms of an average over an effective single-agent process, since the dependence on $\widehat{\mathsf{C}}$, $\widehat{\mathsf{K}}$ and $\widehat{\mathsf{L}}$ can be eliminated using (4.21). To this aim, we now calculate Φ explicitly, in order to compute Eq. (4.21) subsequently. The integration over the variables x_t yields clearly $\prod_t\delta(\widehat{x}_t + w_t)$, so that the integration over $\widehat{x}_t$'s is straightforward. Successively one may perform the Gaussian integration over w_t's:

$$\Phi = \alpha\ \log\int e^{\mathrm{i}\sum_t\widehat{w}_t w_t - \frac{1}{2}\sum_{t,t'}(w_t w_{t'} + \widehat{w}_t L_{tt'}\widehat{w}_{t'} - 2w_t K_{tt'}\widehat{w}_{t'} + w_t C_{tt'}w_{t'})}\ Dw\ D\widehat{w} = \tag{4.23}$$
$$= \alpha\ \log\int e^{-\frac{1}{2}\sum_{t,t'}\widehat{w}_t L_{tt'}\widehat{w}_{t'}}\ e^{-\frac{1}{2}\sum_{t,t'}w_t(\mathsf{E}+\mathsf{C})_{tt'}w_{t'} + i\sum_{t,t'}w_t(\mathsf{I}-\mathrm{i}\mathsf{K})_{tt'}\widehat{w}_{t'}}\ Dw\ D\widehat{w} =$$
$$= -\frac{\alpha}{2}\log\det(\mathsf{E}+\mathsf{C}) + \alpha\int e^{-\frac{1}{2}\sum_{t,t'}\widehat{w}_t L_{tt'}\widehat{w}_{t'} - \frac{1}{2}\sum_{t,t'}\widehat{w}_t\left[(\mathsf{I}+\mathsf{G})^T(\mathsf{E}+\mathsf{C})(\mathsf{I}+\mathsf{G})\right]_{tt'}\widehat{w}_{t'}}\ D\widehat{w}$$

where E is as usual the matrix with all elements equal to one and we have used the fact that $\mathsf{K} = \mathrm{i}\mathsf{G}$. Setting all $L_{tt'}$ to zero in agreement with (4.20) one obtains, after a Gaussian integration,

$$\lim_{\mathsf{L}\to 0} \Phi = -\alpha \ \mathrm{tr} \log(\mathsf{I} + \mathsf{G}) \tag{4.24}$$

whence from (4.21)

$$\widehat{\mathsf{K}}^T = -\alpha \, (\mathsf{I} + \mathsf{G})^{-1} \qquad \text{and} \qquad \widehat{\mathsf{C}} = 0 \tag{4.25}$$

follow. Deriving Φ with respect $L_{tt'}$ and then setting all $L_{tt'}$ to zero, from (4.21) again one obtains

$$\widehat{\mathsf{L}} = -\frac{1}{2}\mathrm{i}\alpha[(\mathsf{I} + \mathsf{G})^{-1} \, (\mathsf{E} + \mathsf{C}) \, \left(\mathsf{I} + \mathsf{G}^T\right)^{-1}] \tag{4.26}$$

4.3.2. Effective single-agent process. Using the above results for $\widehat{\mathsf{K}}^T$, $\widehat{\mathsf{C}}$, and $\widehat{\mathsf{L}}$, we are finally in the position of re-considering the single-agent measure $M(y)$ appearing in the $\langle \ \rangle_*$ average, which now reads

$$M(y) = p[y(0)] \int e^{-\frac{1}{2}\alpha \sum_{t,t'} \widehat{y}(t)[(\mathsf{I}+\mathsf{G})^{-1}(\mathsf{E}+\mathsf{C})(\mathsf{I}+\mathsf{G}^T)^{-1}]_{tt'}\widehat{y}(t')} \times \\ \times e^{\mathrm{i}\sum_t \widehat{y}(t)\left[y(t+1)-y(t)-\alpha\eta s(t)-\theta(t)+\alpha\sum_{t'}[(\mathsf{I}+\mathsf{G})^{-1}]_{tt'}s(t')\right]} D\widehat{y} \tag{4.27}$$

From this it is evident that $M(y)$ can be derived from the effective non-linear single-agent process

$$y(t+1) = y(t) - \alpha \sum_{t'\le t} (\mathsf{I} + \mathsf{G})^{-1}_{tt'} s(t') + \alpha\eta s(t) + \sqrt{\alpha} z(t) + \theta(t) \tag{4.28}$$

where $z(t)$ is a Gaussian noise with zero mean and temporal correlations given by

$$\langle z(t)z(t')\rangle = \left[(\mathsf{I} + \mathsf{G})^{-1}(\mathsf{E} + \mathsf{C})(\mathsf{I} + \mathsf{G}^T)^{-1}\right]_{tt'} = \sum_{ss'} (\mathsf{I} + \mathsf{G})^{-1}_{ts}(\mathsf{E} + \mathsf{C})_{ss'}(\mathsf{I} + \mathsf{G}^T)^{-1}_{s't'} \tag{4.29}$$

According to 4.20, the matrices C and G appearing here are the noise-averaged single-agent correlation and response functions for the process (4.28), with elements

$$C_{tt'} = \langle s(t)s(t')\rangle \qquad \text{and} \qquad G_{tt'} = \left\langle \frac{\partial s(t)}{\partial \theta(t')} \right\rangle \tag{4.30}$$

respectively, while (as always) I is the identity matrix and E denotes the matrix with all entries equal to one. The link between the Markovian multi-agent system (4.5) and the non-Markovian single-agent process (4.28) is established by the fact that, for $N \to \infty$, $C_{tt'}$ and $G_{tt'}$ become identical to the disorder- and agent-averaged correlation and response functions of (4.5):

$$C_{tt'} = \frac{1}{N} \sum_{1\le i\le N} \left[\langle s_i(t)s_i(t')\rangle_{\text{paths}}\right]_{\boldsymbol{a}} \qquad \text{and} \qquad G_{tt'} = \frac{1}{N} \sum_{1\le i\le N} \left[\frac{\partial \langle s_i(t)\rangle_{\text{paths}}}{\partial \theta_i(t')}\right]_{\boldsymbol{a}} \tag{4.31}$$

Before moving to the study of the stationary states of this dynamics, let us make some remarks.

1. By causality, $G_{tt'} = 0$ (and consequently $(\mathsf{I} + \mathsf{G})^{-1}_{tt'}$) for $t - t' \le 0$.
2. By construction, the effective single-agent measure (4.27) is normalized: $\int M(y)dy = 1$.
3. Following step by step the calculation performed in [22] for the $\eta = 0$ case, it is possible to show that, for general η, one has $\lim_{N\to\infty} [\langle A\rangle_{\text{time}}]_{\boldsymbol{a}} = 0$. Moreover, one can prove that

$$\lim_{N\to\infty} \left[\langle A^2\rangle_{\text{time}}\right]_{\boldsymbol{a}} \equiv \lim_{N\to\infty} \left[\sigma^2_\eta\right]_{\boldsymbol{a}} = \frac{1}{2} \lim_{t\to\infty} \left[(\mathsf{I} + \mathsf{G})^{-1}(\mathsf{E} + \mathsf{C})(\mathsf{I} + \mathsf{G}^T)^{-1}\right]_{tt} \tag{4.32}$$

Unfortunately, we have not been able to derive an explicit formula for the stationary volatility starting from (4.32).

4.4. Stationary state in the ergodic regime

Eqs (4.28-4.30) describe the dynamics of the system exactly in the $N \to \infty$ limit. We introduce now the re-scaled quantity $\widetilde{y}(t) = y(t)/t$ (essentially, the time derivative of y), for which we find

$$\widetilde{y}(t) = -\frac{\alpha}{t}\sum_{t'<t}\left[\sum_{t''}(\mathsf{I}+\mathsf{G})^{-1}_{t't''}s(t'') - \eta s(t')\right] + \frac{\sqrt{\alpha}}{t}\sum_{t'<t} z(t') + \frac{1}{t}\widetilde{y}(1) \tag{4.33}$$

We now move to the stationary states of (4.28) upon making the following assumptions:

$$\text{Time-translation invariance (TTI)} \qquad \begin{cases} \lim_{t\to\infty} C_{t+\tau,t} = C(\tau) \\ \lim_{t\to\infty} G_{t+\tau,t} = G(\tau) \end{cases} \tag{4.34}$$

$$\text{Finite integrated response (FIR)} \qquad \lim_{t\to\infty}\sum_{t'\le t} G_{tt'} = \chi < \infty \tag{4.35}$$

$$\text{Weak long-term memory (WLTM)} \qquad \lim_{t\to\infty} G_{tt'} = 0 \quad \forall\, t' \text{ finite} \tag{4.36}$$

These three conditions together ensure that also the retarded self-interaction $\mathsf{R} = (\mathsf{I}+\mathsf{G})^{-1}$ is time-translation invariant. On the basis of other studies on the dynamics of disordered systems, we expect, with these Ansätze, to be able to calculate the dynamics in the stationary state. For long times, i.e. for the variable $\widetilde{y} = \lim_{t\to\infty}\widetilde{y}(t)$, one has

$$\widetilde{y} = -\frac{\alpha s}{1+\chi} + \alpha\eta s + \sqrt{\alpha} z + \theta \tag{4.37}$$

where $s = \lim_{\tau\to\infty}\tau^{-1}\sum_{t<\tau}\mathrm{sgn}[y(t)]$ and $z = \lim_{\tau\to\infty}\tau^{-1}\sum_{t<\tau} z(t)$, while θ is a now static field. The variance of the zero-average Gaussian random variable z can be calculated from (4.29), yielding

$$\langle z^2\rangle = \lim_{\tau,\tau'\to\infty}\sum_{t\le\tau}\sum_{t'\le\tau'}\left[(\mathsf{I}+\mathsf{G})^{-1}(\mathsf{E}+\mathsf{C})(\mathsf{I}+\mathsf{G}^T)^{-1}\right]_{tt'} = \frac{1+c}{(1+\chi)^2} \tag{4.38}$$

with the persistent correlation $c \equiv \langle s^2\rangle = \lim_{\tau\to\infty}\tau^{-1}\sum_{t<\tau} C(t)$. The effective agent is 'frozen' if $\widetilde{y} \neq 0$, so that $s = \mathrm{sgn}(\widetilde{y})$ and he is always employing the same strategy. Setting $\theta = 0$, this is easily seen to be the case if $|z| > \gamma$ with $\gamma = \sqrt{\alpha}\left[(1+\chi)^{-1} - \eta\right]$, provided $\gamma \ge 0$. He is instead fickle when $\widetilde{y} = 0$ or $|z| < \gamma$, and in this case $s = z/\gamma$. A self-consistent equation for c can now be derived by separating the contribution of the frozen agents from that of the fickle ones. Upon defining $\lambda = \gamma/\sqrt{\langle z^2\rangle}$ one finds, after performing the Gaussian integration over z with variance (4.38),

$$c = \langle\Theta(|z|-\gamma)\rangle + \left\langle \Theta(\gamma-|z|)\frac{z^2}{\gamma^2}\right\rangle = \phi + \frac{1}{\lambda^2}\left[\overline{\phi} - \lambda\sqrt{\frac{2}{\pi}}e^{-\frac{\lambda^2}{2}}\right] \tag{4.39}$$

where Θ is the step function, $\overline{\phi} = \mathrm{erf}(\lambda/\sqrt{2})$ is the fraction of fickle agents, and $\phi = 1 - \overline{\phi}$ is the fraction of frozen agents. This equation is identical to that which can be obtained for the self-overlap Q in the replica approach. For $\chi = \left\langle\frac{\partial s}{\partial\theta}\right\rangle = \alpha^{-1/2}\left\langle\frac{\partial s}{\partial z}\right\rangle$ one obtains

$$\chi = \frac{1}{\sqrt{\alpha}}\left\langle\Theta(|z|-\gamma)2\delta(\sqrt{\alpha}z)\right\rangle + \frac{1}{\gamma\sqrt{\alpha}}\left\langle\Theta(\gamma-|z|)\right\rangle = \frac{\overline{\phi}}{\gamma\sqrt{\alpha}} \tag{4.40}$$

Equations (4.38-4.40) form a closed set from which one can solve for ϕ, c and χ for any α and η. Results for c are shown in Figure 4.1. Let us discuss them in detail.

- For negative η, one observes an excellent agreement between theory and experiment for all values of α, implying that none of our assumptions is ever violated.
- When $\eta = 0$, we recover the results of [18,22], which match the simulations perfectly for α larger than the critical value $\alpha_c \simeq 0.3374$. At this point the integrated response χ diverges (FIR is violated) and a transition to a highly non-ergodic regime takes place, where the stationary state depends on the initial conditions $y(0)$. Therefore, in this case, ergodicity breaking is accompanied by a diverging susceptibility and, quite likely, by a breakdown of weak long-term memory as well. This situation is not usual for disordered systems, but it's standard (e.g. ferromagnets). Starting with $y(0) \simeq 0$ leads to a high volatility state, while starting

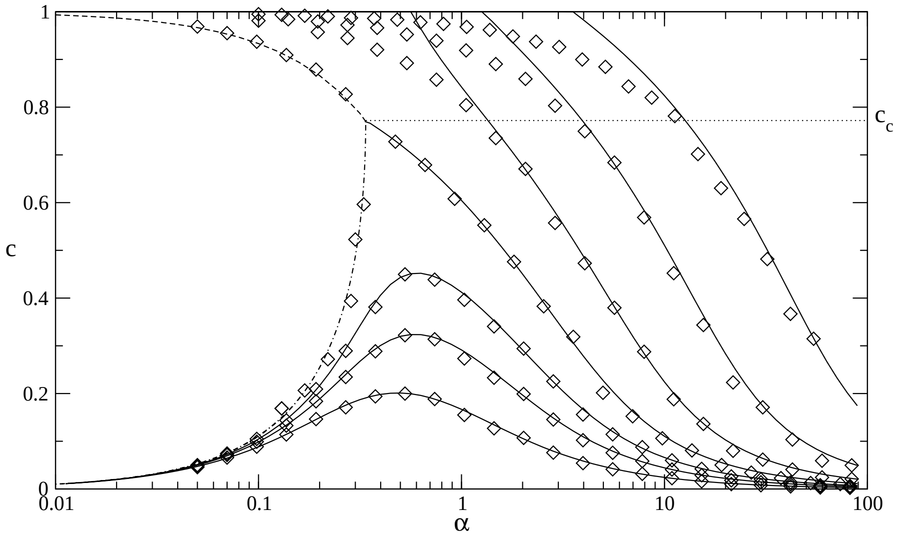

FIGURE 4.1. Persistent correlation c as a function of α for different values of η. Lines represent theoretical predictions. Solid lines: from bottom to top, $\eta = -1, -0.5, -0.25, 0, 0.25, 0.5, 0.7$. Dashed line: $y(0) \gg 0, \eta = 0$. Dot-dashed line: $\eta = 0^-$. Diamonds correspond to computer simulations with $\alpha N^2 = 10,000$, run for 500 time steps and averaged over 50 disorder samples.

with $|y(0)| \gg 1$ leads to relatively low volatility. The latter regime can be solved using the assumption that χ remains very large for all $\alpha < \alpha_c$. In fact, if $\chi \gg 1$ then $\gamma \simeq \sqrt{\alpha}/\chi$ so that $\overline{\phi} = \alpha$, which is equivalent to $\mathrm{erf}(\lambda/\sqrt{2}) = \alpha$. Solving this for λ and inserting the resulting value in (4.39), we obtain the top left branch of the $\eta = 0$ curve in Fig. 4.1, which is again in excellent agreement with numerical results. Finally, as a byproduct, we can also gain some more insight on a result obtained in [25] (see Figure 5 there). The self-overlap Q, static analogous to the persistent autocorrelation c, was numerically found to be linear in α for $\alpha < \alpha_c$ and large initial conditions: $Q(\alpha, y_0 = \infty) \simeq 1 - k\alpha$ with constant k. We are in the position of verifying this conjecture. Starting from (4.39) $\phi = 1 - \alpha$ and $\overline{\phi} = \alpha$, and using $\lambda = \sqrt{2}\mathrm{erf}^{-1}(\alpha)$ we obtain a closed (though complicated) expression for $c(\alpha)$, which can be expanded in powers of α around $\alpha = 0$:

$$c(\alpha) = 1 - \frac{2}{3}\alpha - \frac{\pi}{45}\alpha^3 - \frac{\pi^2}{315}\alpha^5 - \frac{23\pi^3}{37800}\alpha^7 + \mathcal{O}(\alpha^8) \tag{4.41}$$

The correction to the linear term is of order 10^{-4} or less below α_c.

- For positive η, one sees that when $c > c_c \simeq 0.77$ our theoretical predictions deviate from the experimental observations, whereas the agreement is perfect for $c < c_c$. Finding no violation of FIR, we have to conclude that either TTI or WLTM is violated. However, we have found no evidence of aging. Therefore we expect the deviations to be related to the breakdown of WLTM only. This suggest a non-standard picture.

4.5. Ergodicity breaking and the memory-onset condition

4.5.1. Memory-onset condition. The situation in the non-ergodic regime is quite special. Indeed, it is known that not all systems will satisfy the WLTM assumption, but it is still not clear how to treat them analytically (see [4], footnote m on page 188). We here argue as follows. In order to find the onset of memory, we split $G_{tt'}$ in its TTI part and its non-TTI part:

$$\lim_{t\to\infty} G_{tt'} = \widetilde{G}(t-t') + \widehat{G}(t,t') \tag{4.42}$$

We now argue as follows. During the initial stages of the game, small perturbations can cause some agents, which would otherwise have remained fickle, to freeze and vice versa, thus creating a persistent part $\widehat{\mathsf{G}}$ in the response function. As the agents freeze, their state (and consequently their contribution to G) becomes independent of t, so that we expect $\widehat{G}(t,t') = \widehat{G}(t')$. After an initial equilibration period, for all frozen agents the difference between the strategy valuations have become very large, so they are virtually insensitive to perturbations. Hence we must assume that $\lim_{t'\to\infty} \widehat{G}(t') = 0$. The fickle agents, however, remain sensitive to small perturbations. The effects will wear out over time (finite response) and are given by $\widetilde{\mathsf{G}}$.

Assuming $\widehat{\mathsf{G}}$ is small, we expand $(\mathsf{I}+\mathsf{G})^{-1}$ in powers of $\widehat{\mathsf{G}}$ up to first order:

$$(\mathsf{I}+\mathsf{G})^{-1} = (\mathsf{I}+\widetilde{\mathsf{G}})^{-1} - \sum_{n=0}^{\infty}\sum_{m=0}^{n-1}(-\widetilde{\mathsf{G}})^m\widehat{\mathsf{G}}(-\widetilde{\mathsf{G}})^{n-m-1} + \mathcal{O}(\widehat{\mathsf{G}}^2) \tag{4.43}$$

(Note: the time-translation invariant parts of C and G commute, i.e. $[\widetilde{G}\widetilde{C}](\tau) = \int \widetilde{G}(\tau-t)\widetilde{C}(t)dt = \int \widetilde{C}(\tau-t')\widetilde{G}(t')dt' = [\widetilde{C}\widetilde{G}](\tau)$. This is in general not true for the non-TTI parts $\widehat{C}$ and $\widehat{G}$. Hence the power series (4.43). If $\widehat{C}$ and $\widehat{G}$ did commute, one would have obtained $(\mathsf{I}+\mathsf{G})^{-1} = (\mathsf{I}+\widetilde{\mathsf{G}})^{-1} - (\mathsf{I}+\widetilde{\mathsf{G}})^{-2}\widehat{\mathsf{G}} + \mathcal{O}(\widehat{\mathsf{G}}^2)$. The reader may verify that, using the latter expression as a starting point, the final result does not change (up to first order), though the intermediate steps would have been different.) Defining $\widetilde{\chi} = \sum_t \widetilde{G}(t)$ and $\widehat{\chi} = \sum_t \widehat{G}(t)$, one then finds asymptotically

$$\widetilde{y} = -\alpha\left(\frac{1}{1+\widetilde{\chi}} - \eta\right)s + \sqrt{\alpha}z + \alpha\sum_{n=0}^{\infty}\sum_{m=0}^{n-1}(-\widetilde{\chi})^m\sum_{t'}\widehat{G}(t')\sum_{t''}\left[(-\widetilde{\mathsf{G}})^{n-m-1}\right]_{t't''}s(t'') \tag{4.44}$$

Using the rectified linear function $f(x) = x$ for $|x| \le 1$ and $\mathrm{sgn}(x)$ otherwise, we see that if $1/(1+\widetilde{\chi}) > \eta$ then

$$s = f\left(\frac{1}{\widetilde{\gamma}}\left[z + \sqrt{\alpha}\sum_{n=0}^{\infty}\sum_{m=0}^{n-1}(-\widetilde{\chi})^m\sum_{t'}\widehat{G}(t')\sum_{t''}\left[(-\widetilde{\mathsf{G}})^{n-m-1}\right]_{t't''}s(t'')\right]\right) \tag{4.45}$$

where $\widetilde{\gamma} = \sqrt{\alpha}\left(\frac{1}{1+\widetilde{\chi}} - \eta\right)$. As before, we have $\widetilde{\chi} = \alpha^{-1/2}\left\langle\frac{\partial s}{\partial z}\right\rangle$, whereas

$$\widehat{G}(t) = \left\langle\frac{\partial s}{\partial\theta(t)}\right\rangle = \frac{\sqrt{\alpha}}{\widetilde{\gamma}}\sum_{n=0}^{\infty}\sum_{m=0}^{n-1}(-\widetilde{\chi})^m\sum_{t'}\widehat{G}(t')\sum_{t''}\left[(-\widetilde{\mathsf{G}})^{n-m-1}\right]_{t't''}G(t'',t) \tag{4.46}$$

Up to first order in $\widehat{\mathsf{G}}$ one finds, summing over t,

$$\widehat{\chi} = \frac{\widehat{\chi}\widetilde{\chi}\sqrt{\alpha}}{[\widetilde{\gamma}(1+\widetilde{\chi})^2]} + \mathcal{O}(\widehat{G}^2) \tag{4.47}$$

Although $\widehat{\chi} = 0$ is always a solution of this equation (obviously), a bifurcation occurs when

$$\frac{\widetilde{\chi}\sqrt{\alpha}}{\widetilde{\gamma}(1+\widetilde{\chi})^2} = 1 \tag{4.48}$$

This condition possesses the following equivalent formulations:

$$\overline{\phi} = \alpha(1-\sqrt{\eta})^2 \qquad \lambda^2[1+c(\lambda)] = \overline{\phi(\lambda)} \qquad \widetilde{\chi} = \eta^{-1/2} - 1 \tag{4.49}$$

We call this line in the (α,η) plane the memory-onset (MO) line.

4.5.2. Equivalence between the MO-line and the AT-line. We will now show that the MO-line coincides with the AT-line derived in statics. It was found that the replica-symmetric solution becomes unstable when (see (3.54))

$$\lim_{\beta\to\infty}\beta^2\ E_z\left[\left(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2\right)^2\right] = \alpha(1+\chi)^2 \tag{4.50}$$

where $E_z(\ )$ denotes a Gaussian average with zero mean and unit variance, and $\langle\ \rangle_V$ is defined in (3.29). Absorbed a factor $\sqrt{\alpha}$ in β, we can re-write this as

$$\lim_{\beta\to\infty}\beta^2\ E_z\left[\left(\langle\phi^2\rangle_V - \langle\phi\rangle_V^2\right)^2\right] = (1+\chi)^2 \tag{4.51}$$

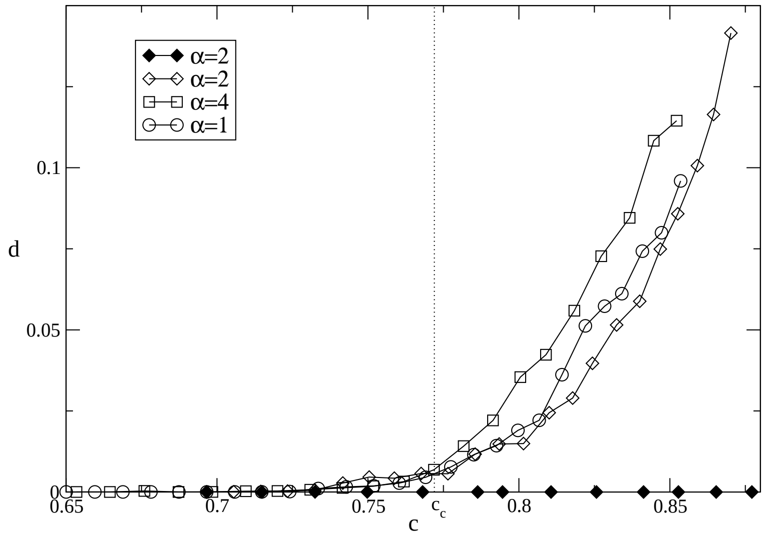

FIGURE 4.2. Distance between the stationary states of two identical copies of the system (see text) as a function of their persistent autocorrelation. Simulations are for various levels of η with $N = 450$, averaged over 100 samples. Open markers correspond to a perturbation at $t = 0$ (for values of α as shown in the legend). Closed markers correspond to a perturbation at $t = 500$ (for $\alpha = 2$). All simulations are run up to 500 steps after the perturbation occurs. Time averages are over the last 300 steps.

Using 3.25 and the replica-symmetric saddle-point equations, together with our definition of γ, it is simple to see that we can set $V_z(\phi) = \frac{1}{2}\gamma\phi^2 - z\phi$, provided we perform the Gaussian average over a new random variable z, with zero mean and variance $\langle z^2 \rangle_z = (1+q)/(1+\chi)^2$, q being the overlap between two different replicas (off-diagonal overlap matrix element in the replica-symmetric setting). If we now define $F(z) = -\lim_{\beta\to\infty} \beta^{-1} \log \int_{-1}^{1} e^{-\beta V_z(\phi)} d\phi$, the AT-line can be written as $\langle F''(z)^2 \rangle_z = (1+\chi)^2$. By Laplace's method we find $F(z) = V_z(\phi^*)$, ϕ^* being the minimum of V_z in $[-1,1]$. For $|z/\gamma| < 1$, ϕ^* lies inside this interval and $V_z(\phi^*) = z^2/(2\gamma)$, while for $|z/\gamma| > 1$ ϕ^* is on the border and $V_z(\phi^*) = \gamma/2 - |z|$. This gives second derivatives that are $-\gamma^{-1}$ and 0, respectively. The AT-line is therefore $E_z[\gamma^{-2}\Theta(1-|z/\gamma|)] + E_z[0\ \Theta(|z/\gamma|-1)] = (1+\chi^2)$, i.e.

$$\langle \gamma^{-2}\Theta(1-|z/\gamma|)\rangle_z = (1+\chi^2) \tag{4.52}$$

Recognizing the term on the l.h.s. as the fraction $\overline{\phi}$ of fickle agents, we find

$$\alpha[1-\eta(1+\chi)^2] = \overline{\phi} \tag{4.53}$$

Introducing λ this is found to be equivalent to $\lambda^2[1+c] = \overline{\phi}$ which is precisely the MO-line.

4.5.3. Discussion and further tests. The MO-line (4.49) further implies that the bifurcation occurs at $c_c \simeq 0.7722$ for $\eta > 0$. Above this value, WLTM can be broken, and indeed one sees from the Figure that numerical results deviate from our theoretical predictions for $c > c_c$. To give further evidence of memory, we have analyzed the time evolution of two identical copies a and b of the system, starting from slightly different initial conditions. We plotted in Figure 4.2 the distance d of the stationary states, given by $(1/N)\sum_i (s_i^a - s_i^b)^2$, where s_i^m is the long-time average of $\mathrm{sgn}(y_i^m)$ $(m = a, b)$, versus the persistent autocorrelation of copy a, c^a. As c^a approaches c_c, the two copies end up in different stationary states, proving that they remember initial conditions[§]. At the same time, if a perturbation is applied much later during the run the copies end up in the same stationary state, indicating that indeed $\widehat{G}(t') \to 0$ as $t' \to \infty$.

In summary: we have studied the long-time limit of the effective single-agent process (4.28) upon assuming time-translation invariance (TTI, (4.34)), finite integrated response (FIR, (4.35)) and weak

[§]The slight bump that occurs before c_c is likely due to the fact that in our simulation the perturbation can not be infinitesimal, but is at least $1/N$.

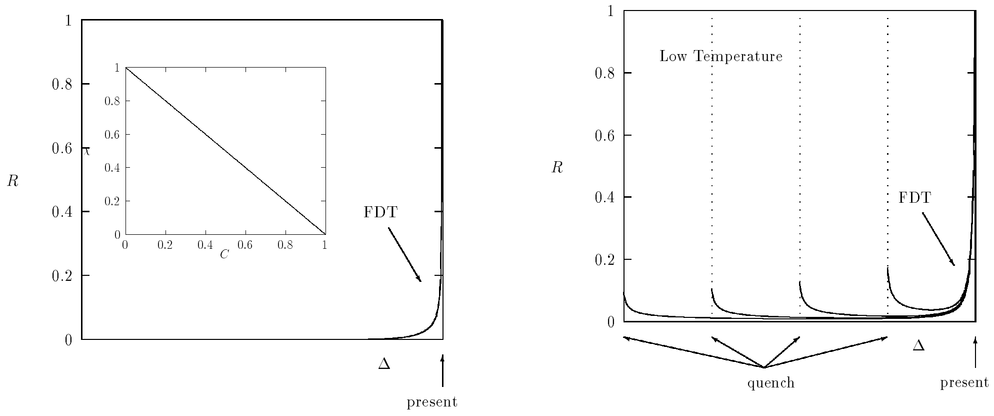

(a) Ergodic regime (b) Non-ergodic regime

FIGURE 4.3. Behaviour of the response function (R here) in mean-field models of spin glasses at (a) high temperatures and (b) low temperatures. From [4].

long-term memory (WLTM, (4.36)) to hold, obtaining an exact and complete description of the persistent order parameters of the 'batch' minority game. For the $\eta = 0$ model, we find that the three assumptions above are perfectly met for $\alpha > \alpha_c \simeq 0.3374$, where the dynamics is ergodic. At α_c the static susceptibility χ diverges, i.e. FIR is violated, and the system undergoes a transition to a highly non ergodic regime ($\alpha < \alpha_c$), where the system's behaviour depends strongly on initial conditions. This suggests that WLTM ceases to hold as well. These two conditions may seem very similar, and thus one would expect them to hold or break down together. However, for $\eta > 0$ we find that WLTM may fail while FIR still holds, implying a continuous onset of long-term memory together with a finite static susceptibility. In particular, this has been shown to happen whenever the persistent autocorrelation $c \equiv \langle s^2 \rangle = \lim_{\tau\to\infty} \tau^{-1} \sum_{t<\tau} C(t)$ becomes larger than the critical value $c_c \simeq 0.772$. The memory-onset condition $c = c_c$ defines a line in the (α, η) plane, which turns out to be identical to the AT-line found in statics. Finding no violation of TTI, we conclude that in this model replica-symmetry breaking (or broken ergodicity) is connected to a continuous onset of long-term memory (breakdown of WLTM) within a TTI regime.

The situation in this model is well explained by the behaviour of the response function. This is a crucial point. For $\eta > 0$ we find no evidence whatsoever of a divergent susceptibility, but the system has acquired a memory:

$$\lim_{t\to\infty} \sum_{t' \le t_w} G_{tt'} > 0 \quad \text{for any } t_w \text{ finite} \tag{4.54}$$

Strong dependence on initial conditions is the only possibility in which the local susceptibility can remain finite. Hence we can compare this situation with that occurring usually in disordered systems (Fig. 4.3). In spin-glasses, the ergodic regime is characterized by a great response to perturbations in the immediate past (see (a)), but no memory of initial conditions. When ergodicity is broken, the response function has a tail stretching down to the quench time. The height of the tail is however decreasing with increasing waiting time, i.e. the initial transients tend to be forgotten (this is the essence of WLTM, see (b)). In Fig. 4.4 one sees the situation in the minority game. The sensitivity to initial conditions is enormous, while the system is insensitive to a perturbation occurred during the time evolution. Just as in the spin-glass case, the local susceptibility remains finite.

There is a simple explanation for this picture, which is however not common in disordered systems. If $\eta > 0$, by lowering α one transforms a dynamical system with one attractor to one with exponentially

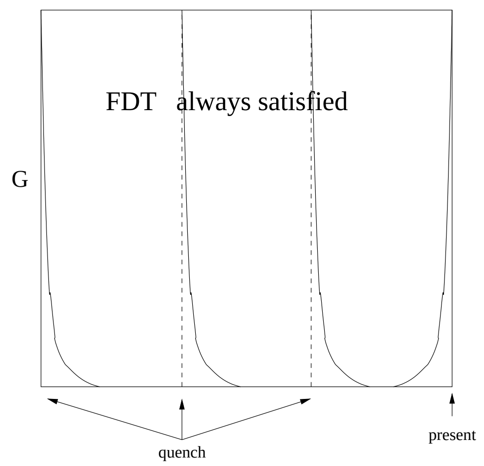

FIGURE 4.4. Sketch of the behaviour of the response function in the minority game in the non-ergodic regime.

many attractors. Let us define $q = (1/N)\sum_{1\le i\le N}\phi_i^a\phi_i^b$ $(a\neq b)$ and $Q = (1/N)\sum_{1\le i\le N}(\phi_i^a)^2$, where a and b label two different replicas of the system, and ϕ_i^a denotes the stationary value of $\mathrm{sign}[y_i]$ in copy a. When only one attractor is present, one expects that $Q = q$. When more attractors are present, the two replicas can either end up in the same state $(Q = q)$, or in two different states $(Q > q)$. This case can only be described by a replica-symmetry breaking formalism: the analogous of the overlap distribution will contain a continuous part.

CHAPTER 5

Variations on the theme

In this final chapter we will describe in broad lines two economics-inspired models based on the original minority game ($\eta = 0$). In the first, we analyze the stationary states of a system of heterogeneous agents competing in a "majority game". This model aims at describing the collective behaviour of a class of economic agents known as "trend-followers", which are believed to be the responsible of market crashes. With the second model, motivated by previous theoretical studies on the fish market, we point at characterizing how efficiently the goods flow in a market of adaptive consumers.

5.1. Economics-inspired problems: agent-based models

The study of markets poses several problems of potential interest to statistical physics in particular. Apart from the interest in studying the stock markets as a genuine complex system, another reason for the increasing popularity of this field among scientists is the joint availability of large historical databases and powerful computers. On the analytic side, the availability of proper mathematical techniques has encouraged the design of agent-based models.

There are at present essentially two broad classes of problems in which the minority game has played an inspiring role. The first one deals with the modeling of financial time series, in particular to retrieve some peculiar statistical regularities known as "stylized facts". Several works have dealt with this problem starting from the minority game setup, and we refer the interested reader to the existing literature [44–46].

The second class of problems deals instead with understanding the macroscopic features of markets from the microscopic dynamical equations governing the behaviour of individual economic agents. Economists classify traders (short term investors) in three main categories, called *fundamentalists*, *chartists* (or *trend-followers*), and *noise traders*, respectively. Fundamentalists believe that the value of a stock is determined by "fundamental" financial facts such as cash flow, earnings, sales, market share, debt levels, etc. Consequently, they expect the price to follow its (slowly changing) fundamental value. Loosely speaking, fundamentalists believe that the market is in a stationary state. Chartists, instead, analyze the charts looking for repetitive patterns in the price dynamics, trying to grasp a trend and extrapolate the tendency for future times. Typically, they follow the moving average over certain horizons of time and ignore the fundamental aspects of the particular stock. Noise traders, finally, are supposed to buy or sell roughly at random in the market. They represent most of the "small" traders which do not follow any reference value and do not look at charts.

The effects of the presence of these three types of traders on the dynamics of prices have been widely discussed in the economic literature. The situation can be summarized as follows. Fundamentalists are believed to pull the price toward its fundamental value, acting as a kind of elastic force. Noise traders introduce a random tendency to deviate from the fundamental value. Chartists destabilize the price causing either a buy or a sell rush. A full theoretical approach is nevertheless still lacking.

5.2. Majority game

5.2.1. Motivation. We will now show that the two main classes of economic agents, namely fundamentalists and chartists (neglecting noise traders), can be mapped onto agents involved in a minority game and in a majority game, respectively. Consider a market in which N agents decide at each time step $t = 1, 2, \ldots$ to invest a certain amount $a_i(t)$ $(1 \leq i \leq N)$ of capital in an asset. We assume that $a_i(t) > 0$ means that agent i will invest $a_i(t)$ Euros (say) to buy stocks, while $a_i(t) < 0$ means that he is wishing to sell a quantity of stocks worth $|a_i(t)|$ Euros. Let $q_i(t)$ be such a quantity. We assume that the price $p(t)$ of the stock at each time step is fixed by the market clearing condition

$S(t)p(t) = D(t)$, $S(t)$ and $D(t)$ being, respectively, the aggregate supply and demand at time t. (Notice that $S(t)$ and $D(t)$ are expressed in different units.) Since the price at time t is fixed only after $S(t)$ and $D(t)$ are known, each $q_i(t)$ will be determined by the last price of the stock, $p(t-1)$, so that $q_i(t) = -a_i(t)/p(t-1)$. Let us define the quantities

$$A^{\pm}(t) = \pm \sum_{1\le i\le N} a_i(t)\Theta[\pm a_i(t)] \qquad \text{and} \qquad A(t) = A^+(t) - A^-(t) = \sum_{1\le i\le N} a_i(t) \tag{5.1}$$

representing, respectively, the total amount of capital agents are willing to invest to buy stocks at time t, the total value of the stocks agents are wishing to sell at time t, and the difference between the former and the latter at time t. It is simple to understand that $D(t)$ is equal to $A^+(t)$, since all agents wishing to buy stocks contribute $a_i(t)$ Euros to the total demand. The total supply, instead, is determined by selling agents ($a_i(t) < 0$) and can be written as $S(t) = \sum_{1\le i\le N} q_i(t)\Theta[-a_i(t)] = A^-(t)/p(t-1)$. The market clearing condition takes the form

$$\frac{p(t)}{p(t-1)} = \frac{A^+(t)}{A^-(t)} \tag{5.2}$$

Roughly speaking, one can say agents who buy today (at price $p(t)$) make a good deal if the price tomorrow ($p(t+1)$) is higher. Vice versa, agents who sell today (at price $p(t-1)$) make a good deal if the price tomorrow ($p(t+1)$) is smaller. Keeping this is mind it is simple to understand that the utility levels (gains) faced by buyers and sellers at time t are respectively given by

$$\begin{aligned} &\text{buyers } (a_i(t) > 0): \qquad u_i(t) = a_i(t)\left[\frac{p(t+1)}{p(t)} - 1\right] \\ &\text{sellers } (a_i(t) < 0): \qquad u_i(t) = a_i(t)\left[\frac{p(t+1)}{p(t-1)} - 1\right] \end{aligned} \tag{5.3}$$

(For buyers, $u_i(t) > 0$ if $p(t+1) > p(t)$; for sellers, $u_i(t) > 0$ when $p(t+1) < p(t-1)$.)

The price tomorrow, $p(t+1)$, is nevertheless unknown to agents today. In order to evaluate their gains, it is then reasonable for them to replace it by an expectation $E_i[p(t+1)]$, which we express with the formula

$$E_i[p(t+1)] = \psi_i p(t) + (1-\psi_i)p(t-1). \tag{5.4}$$

Agents for which $\psi_i < 1$ perform a weighted average of the two previous prices to obtain the expectation of the next price. In this sense, they are somehow assuming that the market is at a stationary state, thus behaving, in economic terms, as *fundamentalists*. Agents for which $\psi_i > 1$, instead, extrapolate a trend from the two previous prices in order to get an expectation for the next one. In other words, they act as *trend followers*. Their time horizon in this case is quite short (just two time steps), however we expect this simple choice to cause no loss of generality. Substituting (5.4) in (5.3) one sees that agents will expect to face a gain

$$\begin{aligned} &\text{buyers } (a_i(t) > 0): \qquad E_i[u_i(t)] = (1-\psi_i)a_i(t)\left[\frac{p(t-1)}{p(t)} - 1\right] \\ &\text{sellers } (a_i(t) < 0): \qquad E[u_i(t)] = -(1-\psi_i)a_i(t)\left[\frac{p(t-1)}{p(t)} - 1\right] \end{aligned} \tag{5.5}$$

Using the market-clearing condition (5.2) this finally becomes

$$\begin{aligned} &\text{buyers } (a_i(t) > 0): \qquad E_i[u_i(t)] = -\frac{1-\psi_i}{A^+(t)}a_i(t)A(t) \\ &\text{sellers } (a_i(t) < 0): \qquad E_i[u_i(t)] = -\frac{1-\psi_i}{A^-(t)}a_i(t)A(t) \end{aligned} \tag{5.6}$$

For $\psi_i < 1$ (i.e. for fundamentalist agents) the expected utilities of both buyers and sellers are qualitatively identical (except for a positive multiplicative factor) to the utility of a minority game player, because the payoff turns out to be positive if they acted contrary to the majority, that is if

$a_i(t)A(t) < 0$. For $\psi_i > 1$ (hence for chartists), instead, the expected utility of both buyers and sellers is positive if they act according to the majority, that is if $a_i(t)A(t) > 0$. We say that they are playing a *majority game*.

One could naively think that the majority game is in some sense trivial. However, as shown below, this is not at all the case, because the limited number of strategies put a strong constraint on the agents' ability to follow the majority.

As said above, trend followers are believed to destabilize markets and cause crashes. Here we will briefly describe the majority game in order to provide a theoretical description of the effects of trend followers on market behaviour. A more comprehensive model of a market can be obtained if one studies a mixed population of fundamentalists and chartists.

5.2.2. Definition of the model. We define here the majority game with heterogeneous agents. The game is played by N agents, and ultimately the limit $N \to \infty$ will be taken. Agents have access to an extensive number of external information patterns, labeled by $\mu \in \{1, \ldots, \alpha N\}$. At each stage of the game $n = 1, 2, \ldots$ an information pattern $\mu(n)$ is chosen at random from $\{1, \ldots, \alpha N\}$ with uniform probabilities and reported to players. All agents receive the same information $\mu(n)$. Each agent is endowed with S different strategies (labeled by $g \in \{1, \ldots, S\}$) to convert the acquired information into a trading decision. Strategies are described by αN-dimensional vectors:

$$\boldsymbol{a}_{ig} = \{a_{ig}^{\mu}\}_{1 \le \mu \le \alpha N} \in \{-1, 1\}^{\alpha N} \qquad (1 \le i \le N \; ; \; 1 \le g \le S) \tag{5.7}$$

By assumption, each a_{ig}^{μ} is selected randomly and independently from $\{-1, 1\}$ with joint probability distribution

$$P(a_{i1}^{\mu}, a_{i2}^{\mu}) = \frac{c}{2}\left(\delta_{a_{i1}^{\mu},1}\delta_{a_{i2}^{\mu},1} + \delta_{a_{i1}^{\mu},-1}\delta_{a_{i2}^{\mu},-1}\right) + \frac{1-c}{2}\left(\delta_{a_{i1}^{\mu},1}\delta_{a_{i2}^{\mu},-1} + \delta_{a_{i1}^{\mu},-1}\delta_{a_{i2}^{\mu},1}\right) \tag{5.8}$$

with $0 \le c \le 1$. One sees we are allowing for a correlation between strategies. For $c = 0$ the two strategies are completely opposite, i.e. $a_{i1}^{\mu} = -a_{i2}^{\mu}$ for all i and μ, in which case, clearly, $\boldsymbol{\omega}_i = \boldsymbol{\Omega} = \boldsymbol{0}$. For $c = 1$, instead, agents only have one strategy, in which case $\boldsymbol{\xi}^{\mu} = \boldsymbol{0}$ for all μ. The usual uniform case corresponds to $c = 1/2$.

Strategies are given an initial valuation or score $p_{ig}(0)$ which is updated at the end of every round. At the start of round n, given $\mu(n)$, every agent selects the strategy with the highest cumulated score, i.e. $\boldsymbol{a}_{i\widetilde{g}_i(n)}$ with

$$\widetilde{g}_i(n) = \arg\max p_{ig}(n) \tag{5.9}$$

and subsequently makes a bid $b_i(n) \in \{-1, 1\}$ according to the trading decision set by the selected strategy, $b_i(n) = a_{i\widetilde{g}_i(n)}^{\mu(n)}$. The total bid at round n is given by

$$A(n) \equiv A^{\mu(n)} = \sum_{1 \le i \le N} b_i(n) \tag{5.10}$$

Finally, for all i and g all scores are updated according to the reinforcement learning dynamics

$$p_{ig}(n+1) = p_{ig}(n) + \frac{1}{N}\, a_{ig}^{\mu(n)} A(n) \tag{5.11}$$

Notice the change of sign with respect to (2.4) with $\eta = 0$.

Restricting to the case $S = 2$, we introduce the variables $y_i(n) = \frac{1}{2}[p_{i1}(n) - p_{i2}(n)]$, $\boldsymbol{\omega}_i = \frac{1}{2}(\boldsymbol{a}_{i1} + \boldsymbol{a}_{i2})$, and $\boldsymbol{\xi}_i = \frac{1}{2}(\boldsymbol{a}_{i1} - \boldsymbol{a}_{i2})$, which we use to write the evolution of the preferences $y_i(n)$ (analogous of (2.10) in the minority game) as

$$y_i(n+1) = y_i(n) + \frac{1}{N}\, \xi_i^{\mu(n)} A(n) = y_i(n) + \frac{1}{N}\, \xi_i^{\mu(n)} \left[\Omega^{\mu(n)} + \sum_{1 \le j \le N} \xi_j^{\mu(n)} s_j(n)\right] \tag{5.12}$$

where, as usual, $s_i(n) = \mathrm{sign}[y_i(n)]$. We are interested in analyzing the stationary state of this process.

5.2.3. Statics. Following the line of reasoning adopted for the minority game, it is simple to conclude that in the majority game agents minimize the function

$$H(\boldsymbol{\phi}) = -\overline{\langle A \rangle^2} = -\frac{1}{\alpha N} \sum_{\mu \leq \alpha N} \left(\Omega^\mu + \sum_{i \leq N} \phi_i \xi_i^\mu \right)^2 \tag{5.13}$$

where the same notation as in Chapter 2 and 3 has been used. As was reasonable to expect, $H(\boldsymbol{\phi})$ is highly reminiscent of the Hopfield model of neural networks (notice however that in this case the spins $\phi \in [-1, 1]$ are continuous variables and that the "stored patterns" ξ_i^μ are random variables taking on values in $\{-1, 0, 1\}$). Moreover, we still have a zero-temperature process.

Some informations about the stationary states can be extended from the knowledge of the Hopfield model [11]. If the number of information patterns remained finite as $N \to \infty$, we would expect that triggered by correlations between the initial state and one pattern, say $\boldsymbol{\xi}^1$, the state vector $\boldsymbol{\phi}$ evolves towards $\boldsymbol{\xi}^1$ (in the context of neural networks, one says that pattern $\boldsymbol{\xi}^1$ has been recalled or retrieved). The similarity between the state vector and the "stored patterns" is measured by the so-called (normalized) *overlaps* $m_\mu(\boldsymbol{\phi}) = (1/N) \sum_{i \leq N} \xi_i^\mu \phi_i$. Therefore we would expect to find that, in the steady state, $m_1(\boldsymbol{\phi}) \neq 0$ asymptotically whereas $m_\mu(\boldsymbol{\phi}) = \mathcal{O}(N^{-1/2})$ for $\mu \geq 2$. However, we have to deal with extensively many patterns. As a consequence, and on the basis of well-known properties of the Hopfield model with extensively many patterns, we expect to find that, for sufficiently large values of α the interference between patterns will destroy retrieval and asymptotically $m_\mu(\boldsymbol{\phi}) = 0$ for all μ. However, for α smaller than a critical value α_c (unrelated to the one found in the minority game) recall of (at least) the $\mu = 1$ pattern will become possible and $m_1(\boldsymbol{\phi}) \neq 0$. Notice that "retrieval" in this case is a crucial condition. It signals essentially jumps in the price (either upwards or downwards). The fact that majority game players are mapped onto a Hopfield Hamiltonian is then in some sense a confirmation of the fact that trend followers cause either buy or sell rushes. In particular, if $\mu = 1$ labels the macroscopic pattern, every time the randomly chosen information is 1, a finite fraction of the agents will buy (or sell) and the price will go up (or down) steeply.

Notice that for $c = 0$ one has $H(\boldsymbol{\phi}) = -(\alpha N)^{-1} \sum_{\mu \leq \alpha N} (\sum_{1 \leq i \leq N} \xi_i^\mu \phi_i)^2$, and the statics of the Hopfield model should be exactly recovered. For $c = 1$, instead, $H(\boldsymbol{\phi}) = -(\alpha N)^{-1} \sum_{\mu \leq \alpha N} (\sum_{1 \leq i \leq N} \omega_i^\mu)^2$ with $\omega_i^\mu \in \{-1, 1\}$ independent, identically-distributed (with equal probability) random variables.

It is important to say that in our case we are not interested in the minima of H, but, rather, in its local minima. To this aim, we can adapt the calculations done around the Hopfield model [11,47,48], and proceed to the replica analysis of the minima of H, concentrating on

$$[\log Z]_{\boldsymbol{a}} = \lim_{n \to 0} \frac{1}{n} \log \left[Z^n \right]_{\boldsymbol{a}} \qquad Z = \mathrm{Tr}_{\boldsymbol{\phi}} e^{-\beta H_\eta(\boldsymbol{\phi})} \tag{5.14}$$

where, as before, $[\ ]_{\boldsymbol{a}}$ denotes an average over disorder.

After carrying out the disorder average and separating the "recalled" pattern $\mu = 1$ from the rest (as usually done in Hopfield models, see e.g. [47, 48]) we obtain

$$[Z^n]_{\boldsymbol{a}} = \int e^{-\beta n N f(\boldsymbol{z}^1, \mathsf{Q}, \mathsf{R})} \, d\boldsymbol{z}^1 \, d\mathsf{Q} \, d\mathsf{R} \tag{5.15}$$

where Q is the overlap matrix with elements $Q_{ab} = (1/N) \sum_{1 \leq i \leq N} \phi_i^a \phi_i^b$, $\mathsf{R} = \{R_{ab}\}_{1 \leq a,b \leq n}$ is the matrix of conjugate variables of the overlaps, $\boldsymbol{z}^1 = \{z_a^1\}_{1 \leq a \leq n}$, and

$$\begin{aligned} f(\boldsymbol{z}^1, \mathsf{Q}, \mathsf{R}) = \frac{1}{2n} \sum_{a \leq n} \left(z_a^1 \right)^2 + \frac{\alpha}{2\beta n} \log \det \mathsf{T} + \frac{\alpha\beta}{2n} \sum_{1 \leq a,b \leq n} R_{ab} Q_{ab} + \\ - \frac{1}{\beta n} \log \left[\mathrm{Tr}_{\{\phi^1, \dots, \phi^n\}} e^{\beta \sqrt{\frac{2}{\alpha}} \sum_{1 \leq a \leq n} z_a^1 (\omega^1 + \phi^a \xi^1) + \frac{\alpha\beta^2}{2} \sum_{a,b \leq n} R_{ab} \phi^a \phi^b} \right]_{\boldsymbol{a}^1} \end{aligned} \tag{5.16}$$

Here, $\mathsf{T} = \mathsf{I} - (2\beta/\alpha)[c\mathsf{E} + (1 - c)\mathsf{Q}]$, and $[\ ]_{\boldsymbol{a}^1}$ denotes an average over the possible choices of the (random) "recalled" pattern.

Imposing the replica-symmetric Ansatz

$$z_a^1 = z \qquad Q_{ab} = Q\delta_{ab} + q(1-\delta_{ab}) \qquad R_{ab} = R\delta_{ab} + r(1-\delta_{ab}) \tag{5.17}$$

we arrive at the following expression for the replica-symmetric "free energy density" f:

$$f(\boldsymbol{z}^1_{\mathrm{RS}}, \mathsf{Q}_{\mathrm{RS}}, \mathsf{R}_{\mathrm{RS}}) = \frac{1}{2}z^2 + \frac{\alpha}{2\beta}\log\left[1 - \frac{2\beta}{\alpha}(1-c)(Q-q)\right] - \frac{c+(1-c)q}{1-\frac{2\beta}{\alpha}(1-c)(Q-q)} + \\ + \frac{\alpha\beta}{2}(RQ - rq) - \frac{1}{\beta}E_y\left[\log\left[\int_{-1}^{1} e^{-\beta V_y(\phi)}d\phi\right]_{\boldsymbol{a}^1}\right] \tag{5.18}$$

where

$$V_y(\phi) = -z\sqrt{\frac{2}{\alpha}}\left(\omega^1 + \phi\xi^1\right) - y\phi\sqrt{\alpha r} - \frac{\alpha\beta}{2}(R-r)\phi^2. \tag{5.19}$$

We remark that in order to capture the minima of H we have to study the $\beta \to \infty$ limit of the free energy.

The corresponding saddle-point equations read

$$r = \frac{4(1-c)[c+(1-c)q]}{\alpha^2(1-\chi)^2} \qquad \beta(R-r) = \frac{2}{\alpha}\frac{1-c}{1-\chi} \tag{5.20}$$

$$Q = E_y[\langle\phi^2\rangle_V] \qquad \beta(Q-q) = \frac{1}{\sqrt{\alpha r}}E_y[\langle\phi y\rangle_V] \qquad z = \sqrt{\frac{2}{\alpha}}(1-c)\,\langle\phi\rangle_V \tag{5.21}$$

where

$$\chi = \frac{2\beta}{\alpha}(1-c)(Q-q) \qquad \text{and} \qquad \langle\cdots\rangle_V = \frac{\int_{-1}^{1}\cdots\; e^{-\beta V_y(\phi)}d\phi}{\int_{-1}^{1} e^{-\beta V_z(\phi)}d\phi} \tag{5.22}$$

In the limit $\beta \to \infty$ the thermal averages can be computed using a battery of error-functions, yielding

$$\langle\phi^2\rangle_V = 1 \qquad \langle\phi y\rangle_V = \frac{2}{\sqrt{2\pi}}e^{-y_0^2/2} \qquad \langle\phi\rangle_V = \mathrm{erf}(y_0/\sqrt{2}) \tag{5.23}$$

where $y_0 = \sqrt{2/r}(1-c)z/\alpha$. One sees that the saddle-point value of Q is one, meaning that agents in the stationary state are all "frozen" (one has either $\phi_i = 1$ or $\phi = -1$). Moreover, we can identify the quantity $\langle\phi\rangle_V$ with the overlap with the macroscopic pattern. In the $\beta \to \infty$ limit we can look as usual for solutions with $q = Q = 1$ and $\chi < \infty$. The quantity $x = y_0/\sqrt{2}$ is found to satisfy the transcendental self-consistency condition

$$x\sqrt{\frac{2\alpha}{1-c}} = F(x) \qquad F(x) = (1-c)\mathrm{erf}(x) - \frac{2x}{\sqrt{\pi}}e^{-x^2} \tag{5.24}$$

that for $c = 0$ coincides with the result obtained in the Hopfield model [47]. Notice that $F(x)$ is anti-symmetric, implying that solutions of the self-consistency equation (5.24) will come in pairs $(x, -x)$. This fact reflects nothing but the symmetry of the Hopfield Hamiltonian, to which (5.13) is identical for $c = 0$ and at zero temperature, with respect to an overall state-flip $\boldsymbol{\phi} \to -\boldsymbol{\phi}$.

5.2.4. Results. Numerical solution of equation (5.24) leads to the following situation (see Fig. 5.1). For all $0 \le c \le 1$ there is a value $\alpha_c(c)$ of α separating a no-retrieval phase ($\alpha > \alpha_c$) where $x = 0$ (and hence $\langle\phi\rangle_V = 0$) is the only solution, from a retrieval phase ($\alpha < \alpha_c$) where two non zero solutions arise. The transition is first-order, similarly to what happens in the Hopfield model. When $c = 0$ we obviously find $\alpha_c(0) \simeq 0.138\ldots$ again in agreement with the Hopfield model. For $c = 1$ the memory phase disappears. The line of critical points $\alpha_c(c)$ is shown in Fig. 5.1 (a). On this line, the overlap with the macroscopic pattern jumps from a finite value discontinuously down to zero (just like in the Hopfield case), as shown in Fig. 5.1 (b). In the retrieval phase, market crashes or buying rushes occur. Preliminary numerical simulations have revealed a very rich behaviour, comparable to that of the minority game. We defer the reader to forthcoming works for more details on these points [28].

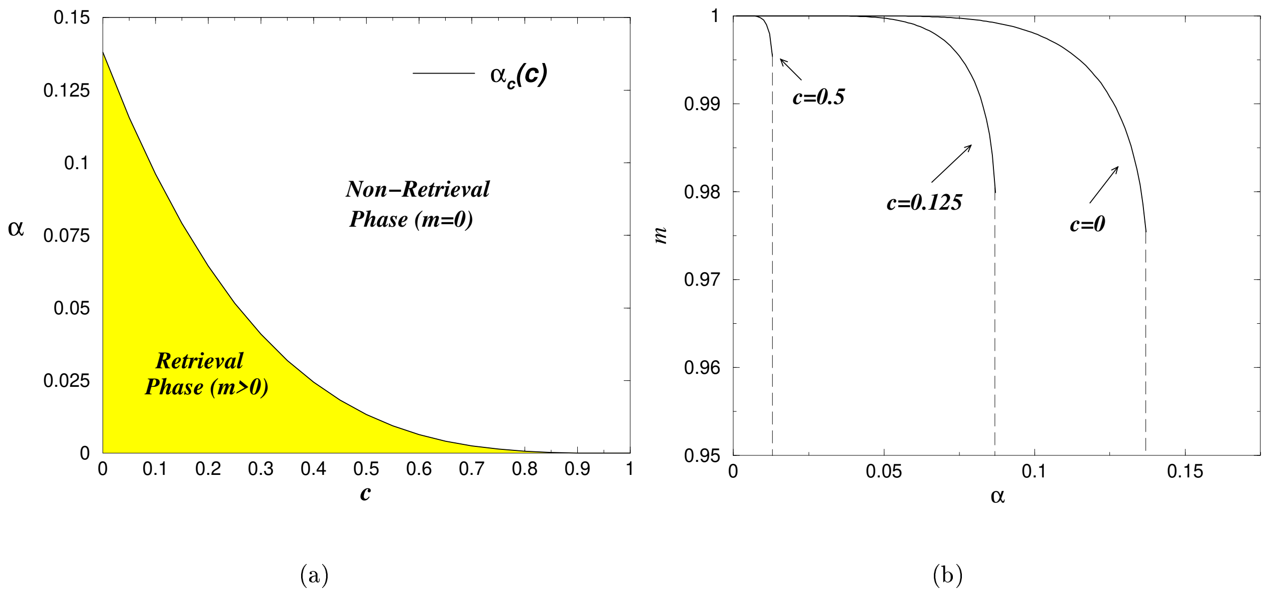

FIGURE 5.1. (a) Phase diagram of the majority game in the (α, c) plane. (b) Behaviour of the normalized overlap with the macroscopic pattern for different values of c.

5.3. A model for the fish market

5.3.1. The problem. At a highly simplified pictorial level, a market, like several other social or economic systems, can be thought of as a network whose nodes are the economic agents (traders, companies, etc.) and whose links or bonds represent the trading relationships existing between them. The evolution of the market can then be visualized as the evolution of the trading links inside the network. Clearly, the structure and time-evolution of the network determines, to a large extent, the global outcome of the market interactions, like the value of macroscopic observables such as prices. From a theoretical viewpoint, one can tackle the problem of understanding the network-dependence of macroscopic quantities from the outset, by assuming a network structure and defining some simple dynamical rules governing the formation or deletion of links. An alternative approach, which to our knowledge started with [49], consists in defining microscopic behavioural (adaptive) rules for the agents and studying the emergent collective behaviour. Within this framework, two questions have been addressed. (i) Under which conditions do stable trading relationships arise? (ii) How does their formation affect the outcome of the network (prices)? By stable links we mean those existing between agents who trade preferentially with each other ("frozen" links).

The formation of stable trading relationships has been experimentally observed in several types of markets, especially in those in which agents trade frequently with each other, like the markets for perishable goods. In this type of context, sellers cannot (in principle) hold inventories, so that it would be convenient for them to bring to the market just the quantity they expect to sell. Clearly, they would be able to predict the needed supply with a much better accuracy should they have fixed customers. In turn, should buyers search too long for the best deal in the market, they would run the risk of not finding what they want. Having a fixed supplier would enormously reduce this possibility. The formation of stable trading links between buyers and sellers is therefore mutually profitable in this case. However, sellers should not take advantage of the situation (e.g. by raising prices), since they could lose customers. Therefore, the price dynamics (including fluctuations) is expected to depend strongly on whether there are stable trading relationships or not.

Although very schematic, the model of [49] provides a qualitative insight on the situation. Agents are assumed to follow a simple learning dynamics, in which they choose a supplier at each time step basing on their past experience. Upon varying the parameters of the model, one finds a transition between a *disordered* phase, where no stable links exist and consumers flip from one supplier to the other with a probability distribution that is almost uniform over suppliers, and an *ordered* phase, in which some producers are visited more frequently than others.

Here we introduce a variation on minority-game's theme to study a closely related problem, namely how efficiently do the goods flow from producers to consumers in a market of heterogeneous agents. In our model, efficiency is reflected in prices and price-fluctuations.

5.3.2. The model. Our model is defined as follows. We consider a market in which a single asset (say, fish) is traded, with N buyers or consumers (labeled by roman indices i, j, etc.) and P sellers or producers (labeled by μ). Each consumer is endowed with S different strategies to buy in the market, representing his possible consumption plans, denoted by M-dimensional vectors

$$\boldsymbol{q}_{ig} = \{q_{ig}^{\mu}\}_{1\le\mu\le P} \in \{0,1\}^P \qquad (i=1,\dots,N; g=1,\dots,S) \tag{5.25}$$

q_{ig}^{μ} should be interpreted as the quantity of goods consumer i demands from seller μ under his g-th consumption plan. We assume that each component q_{ig}^{μ} is selected randomly and independently from $\{0,1\}$ with probability distribution

$$P(q_{ig}^{\mu}) = p\delta_{q_{ig}^{\mu},0} + (1-p)\delta_{q_{ig}^{\mu},1} \qquad 0\le p\le 1 \tag{5.26}$$

for all i, g, and μ, and fixed. In this case, we say that with probability $1-p$ consumer i visits and, eventually, buys one unit of goods from m, while with probability p he does not visit m. Rounds of the market (or time steps) are denoted by $n = 0, 1, \dots$. In the beginning ($n = 0$), each strategy of every agent is given an initial valuation $p_{ig}(0)$, which is updated at the end of every round. At the start of round n, every agent selects the strategy with the highest valuation, $\widetilde{g}_i(n) = \arg\max p_{ig}(n)$, and demands goods from the producers according to the selected consumption plan $\boldsymbol{q}_{i\widetilde{g}_i(n)}$. Consequently, each producer μ receives a total demand $Q^{\mu}(n) = \sum_{i\le N} q_{i\widetilde{g}_i(n)}^{\mu}$ and fixes his price accordingly. In general, a large $Q^{\mu}(n)$ will imply a large price. Prices here are not defined directly but, rather, through the utility experienced by consumers in buying from producer μ, which we denote by $G[Q^{\mu}(n)]$. G is called *marginal utility.* It is a monotonically decreasing function of its argument (the higher the price, the worse the deal), and for simplicity we take $G[Q^{\mu}(n)] = -Q^{\mu}(n)$. Finally, for all i and g all strategy valuations are updated according to a reinforcement learning dynamics of the form

$$p_{ig}(n+1) = p_{ig}(n) + \frac{1}{P}\sum_{\mu\le P} q_{ig}^{\mu} G[Q^{\mu}(n)] = p_{ig}(n) - \frac{1}{P}\sum_{\mu\le P} q_{ig}^{\mu} Q^{\mu}(n) \tag{5.27}$$

and agents move to the next round. Two remarks are now straightforward. Firstly, agents neglect their market impact completely, so this is a kind of $\eta = 0$ minority game; secondly, the structure of the learning dynamics is reminiscent of the batch process.

We will be interested in the $N \to \infty$ limit, assuming that $\alpha = P/N$ remains finite. There are two quantities whose asymptotic behaviour for large systems we study. Denoting by on over-line the average over producers, i.e. $\overline{x^{\mu}} = (1/P)\sum_{\mu\le P} x^{\mu}$, the first one is given by

$$H = \overline{\langle A\rangle_{\text{time}}^2} = \overline{\left[\lim_{L\to\infty}\frac{1}{L-\ell}\sum_{\ell\le n\le L} A^{\mu}(n)\right]^2} \qquad A^{\mu}(n) = Q^{\mu}(n) - \overline{Q^{\mu}(n)} \tag{5.28}$$

where ℓ is the time needed to achieve stationarity. H measures the uniformity of the distribution of marginal utility over producers. When $H > 0$ there is at least one producer which consumers find more convenient than the others. In this case, the good flows inefficiently. When $H = 0$, instead, no producer is more convenient for consumer, in which case the flow is efficient (evenly distributed among producers). Notice that H is related to the average price. The second interesting quantity is the volatility

$$\sigma^2 = \overline{\langle A^2\rangle_{\text{time}}} = \overline{\lim_{L\to\infty}\frac{1}{L}\sum_{\ell\le n\le L}[A^{\mu}(n)]^2} \tag{5.29}$$

In particular, price fluctuations are given by $\sigma^2 - H$. In the Fig. 5.2 (a) we show the behaviour of H, σ^2 and $\sigma^2 - H$ as observed in numerical simulations. One sees that for sufficiently large α the market

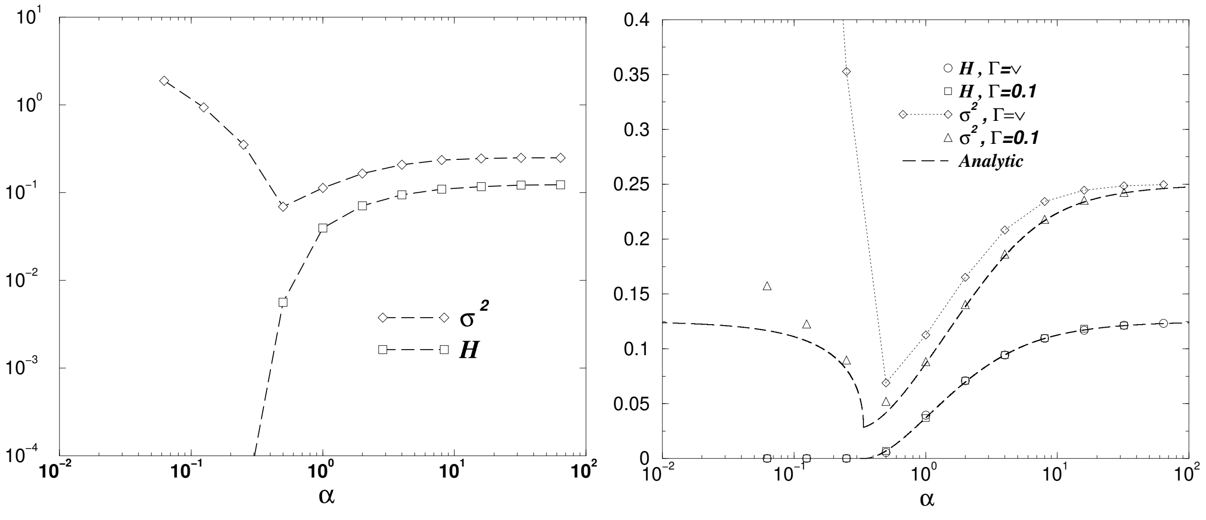

(a) (b)

FIGURE 5.2. (a) Numerical results from simulations of the fish market model, sampled over 200 realizations of the quenched disorder. (b) Comparison between analytic results and simulations for different values of the "learning rate" Γ (see text).

is inefficient ($H > 0$). Correspondingly, however, price fluctuations ($\sigma^2 - H$) are small (though the price is higher on the average). For low α, instead, the market is efficient, and the average price is lower, but fluctuations are much larger. Interestingly enough, the most efficient market occurs at the critical point $\alpha_c \simeq 0.3374$, where $H = 0$ and the fluctuations attain a minimum (see also [46]).

Before turning to the analytic study of the model (which will be outlined next) we make an important remark. Due to the nature of the random variables $\boldsymbol{q}_{ig}$, it is simple to understand that $Q^\mu(n)$ scales with N and p as $\mathcal{O}(\sqrt{Np(1-p)})$. One can therefore argue that in the limit $N \to \infty$ the results of the $\eta = 0$ minority game for the density of macroscopic observables will be recovered, modulo a factor $p(1-p)$. This conclusion turns out to be fully correct.

5.3.3. Results. We will now show that the dynamics of the fish market model minimizes a global macroscopic function. We start by defining

$$u_i^\mu(n) = -q_{ig}^\mu Q^\mu(n) \qquad \text{and} \qquad u_i(n) = \sum_{1\le\mu\le P} u_i^\mu(n) \tag{5.30}$$

Using this, the learning dynamics reads $p_{ig}(n+1) = p_{ig}(n) + \Delta_{ig}(n)$ where $\Delta_{ig}(n) = \frac{1}{P}u_i(n)$. We further introduce the regularized learning dynamics:

$$\text{Prob}\,\{\widetilde{g}_i(n) = g\} \equiv \pi_{ig}(n) = C\exp\left[\Gamma p_{ig}(n)\right] \tag{5.31}$$

Taking the continuous-time limit in the rescaled time $\tau = t/P$ as done e.g. in [18] one obtains

$$\partial_t \pi_{ig}(t) = \Gamma \pi_{ig}(t)\left[\partial_t p_{ig}(t) - \sum_{1\le g\le S} \pi_{ig}(t)\partial_t p_{ig}(t)\right] = \Gamma\pi_{ig}(t)\left[\overline{\langle\Delta_{ig}\rangle} - \sum_{1\le g\le S}\pi_{ig}(t)\overline{\langle\Delta_{ig}\rangle}\right] \tag{5.32}$$

where $\overline{\cdots}$ = average over producers $= (1/P)\sum_{1\le\mu\le P}\cdots$, $\langle\ \rangle$ = time average in the stationary state of (5.31).

In a more compact notation,

$$\partial_t \pi_{ig} = \Gamma\pi_{ig}(t)\left[\overline{\langle\Delta_{ig}\rangle} - \boldsymbol{\pi}_i \cdot \overline{\langle\boldsymbol{\Delta}_i\rangle}\right] \tag{5.33}$$

where $\boldsymbol{x}_i = (x_{ig})_{1\le g\le S}$ and $(\boldsymbol{x}_i, \boldsymbol{y}_i) = \sum_{1\le g\le S} x_{ig}y_{ig}$. Let us compute $\overline{\langle\Delta_{ig}\rangle}$. Recalling the definition of Δ_{ig} and the fact that the only time dependence is hidden in Q^μ, we obtain

$$\langle\Delta_{ig}\rangle = -\frac{1}{P}\sum_{1\le\mu\le P}\left[q_{ig}^\mu \sum_{1\le j\le N} \boldsymbol{\pi}_j\cdot \boldsymbol{q}_j^\mu\right] = -\frac{1}{P}\sum_{1\le\mu\le P}\left[q_{ig}^\mu \sum_{1\le j\le N}\left\langle \boldsymbol{q}_j^\mu\right\rangle\right] = -\overline{q_{ig}^\mu \sum_{1\le j\le N}\left\langle \boldsymbol{q}_j^\mu\right\rangle} \tag{5.34}$$

It is obvious that $\overline{\langle\Delta_{ig}\rangle} = \langle\Delta_{ig}\rangle$. Furthermore,

$$\boldsymbol{\pi}_i\cdot\overline{\langle\boldsymbol{\Delta}_i\rangle} = -\frac{1}{P}\sum_{1\le\mu\le P}\left[(\boldsymbol{\pi}_i\cdot \boldsymbol{q}_i^\mu)\sum_{1\le j\le N}\boldsymbol{\pi}_j\cdot \boldsymbol{q}_j^\mu\right] = -\overline{\langle\boldsymbol{q}_i^\mu\rangle\sum_{1\le j\le N}\left\langle \boldsymbol{q}_j^\mu\right\rangle} \tag{5.35}$$

Putting things together we get

$$\partial_t\pi_{ig}(t) = -\Gamma\pi_{ig}(t)\sum_{1\le j\le N}\left(\overline{q_{ig}^\mu\left\langle \boldsymbol{q}_j^\mu\right\rangle} - \overline{\langle\boldsymbol{q}_i^\mu\rangle\left\langle \boldsymbol{q}_j^\mu\right\rangle}\right) \tag{5.36}$$

It is simple to show that this dynamics minimizes the function

$$\mathcal{H} = \overline{\langle Q\rangle^2} = \overline{\left[\sum_{1\le i\le N}\left(\sum_{1\le g\le S}\pi_{ig}q_{ig}^\mu\right)\right]^2} \tag{5.37}$$

In fact, $\partial_{\pi_{ig}}\mathcal{H} = \overline{2q_{ig}^\mu\langle Q\rangle} = -2\partial_t p_{ig}(t)$, so that

$$\begin{aligned}\partial_t\mathcal{H} &= \sum_{1\le i\le N}\sum_{1\le g\le S}\left(\partial_{\pi_{ig}}\mathcal{H}\right)\left(\partial_t\pi_{ig}(t)\right) = \\ &= -2\Gamma\sum_{1\le i\le N}\sum_{1\le g\le S}\pi_{ig}(t)\left(\partial_t p_{ig}(t)\right)\left[\left(\partial_t p_{ig}(t)\right) - \boldsymbol{\pi}_i(t)\cdot\partial_t\boldsymbol{p}_i(t)\right] = \\ &= -2\Gamma\sum_{1\le i\le N}\sum_{1\le g\le S}\left[\pi_{ig}(t)\left(\partial_t p_{ig}(t)\right)^2 - \pi_{ig}(t)\left(\partial_t p_{ig}(t)\right)\boldsymbol{\pi}_i(t)\cdot\partial_t\boldsymbol{p}_i(t)\right] = \\ &= -2\Gamma\sum_{1\le i\le N}\left[\left\langle\left(\partial_t\boldsymbol{p}_i(t)\right)^2\right\rangle - \left\langle\partial_t\boldsymbol{p}_i(t)\right\rangle^2\right] \le 0\end{aligned} \tag{5.38}$$

The dynamics therefore converges to the local minima of $\mathcal{H}$, which can be studied with replica techniques. It can be shown that for $\Gamma = 0$ (5.37) is exactly a Lyapunov function of the dynamics. For $\Gamma > 0$ this is in general not true, however the stationary states are still well described by the minima of $\mathcal{H}$.

Notice that $\mathcal{H}$ can be thought of as the Hamiltonian of a mean-field spin system with soft spins taking values in $[0, 1]$. The replica calculation for this case is a straightforward generalization of more usual, and we won't report it here. We limit ourselves to a discussion of the results, in comparison to numerical simulations (see Fig. 5.2 (b)). For all α, the minimum of $\mathcal{H}$ provides us with a perfect agreement with the numerical results for H (as one could easily predicted), for all values of Γ. The situation for σ^2/N is slightly more complex. Numerical results for infinite Γ are not described faithfully by the minima of $\mathcal{H}$. This is related to the fact that the static approximation in principle only holds for small Γ. In fact, the agreement between analytic results and computer simulations with $\Gamma = 0.1$ is perfect for all $\alpha > \alpha_c$.

CHAPTER 6

Epilogue

In this thesis we have applied concepts and techniques from equilibrium and non-equilibrium statistical mechanics of disordered systems to analyze economics-inspired problems. The main part of the work was dedicated to the solution of the minority game with random external information and market-impact correction. We have shown how this model maps onto a mean-field spin-glass with non-standard dynamical properties. This behaviour has been compared with known results on the dynamics of glassy systems, and its possible origin has been discussed. The combined use of static and dynamic techniques has allowed us to obtain a quite precise description of the situation, which has been confirmed convincingly by computer simulations. Similar static calculations have been performed for two other models bearing a more direct economic interpretation.

We believe that statistical mechanics techniques will have a lasting impact on economic and financial modeling. In these fields, in fact, the main stumbling block of spin-glass theory, that is the issue of the extension of mean-field results to short-range models, is absent. Mean-field interactions can be properly justified. This state of things is in some sense analogous to that of the theory of neural networks, or of combinatorial optimization problems. In those cases, the spin-glass analogy has been found to be a crucial ingredient to obtain new and deep results. In turn, one often faces classes of systems which are close relatives of mean-field spin-glasses, as for instance diluted spin-glasses for combinatorial optimization problems, and yet require refinements of the available techniques and, in some cases, new physical insight to be solved.

Of course, the modeling of economic systems starts with a high level of schematization, as in the minority game. The possibility to study the collective behaviour of systems of heterogeneous agents interacting through a market-type mechanism represents already a very good and promising starting point. However, market economics offers enough challenging problems and interesting puzzles to keep statistical physicists busy for years to come. The remarkable versatility of our methods, especially for the study of non-equilibrium systems, allows to obtain solutions in the form of macroscopic laws for large systems that need not satisfy dubious or even unacceptable assumptions, as those made by a part of mainstream economic theory. This makes us confident that further progress will be achieved as more complex and realistic models will be explored in the near future. We hope that at that point the work presented in this thesis will continue to be useful in explaining at least at a basic and qualitative level the operation of this fascinating class of systems.

Bibliography

[1] Mézard M, Parisi G, and Virasoro MA. *Spin glass theory and beyond.* World Scientific (Singapore), 1987.
[2] Fischer KH and Hertz JA. *Spin glasses.* University Press (Cambridge, UK), 1991.
[3] Young AP, editor. *Spin glasses and random fields.* World Scientific (Singapore), 1998.
[4] Bouchaud J-P, Cugliandolo LF, Kurchan J, and Mézard M. Out of equilibrium dynamics of spin-glasses and other glassy systems. In [3], pages 161–223. Preprint cond-mat/9702070.
[5] Martin PC, Siggia ED, and Rose HA. Statistical dynamics of classical systems. *Phys. Rev. A*, 8:423–437, 1973.
[6] Zinn-Justin J. *Quantum field theory and critical phenomena.* University Press (Oxford, UK), 3rd edition, 1996.
[7] De Dominicis C. Dynamics as a substitute for replicas in systems with quenched random impurities. *Phys. Rev. B*, 18:4913–4919, 1978.
[8] Cugliandolo LF and Kurchan J. On the out of equilibrium relaxation of the sherrington–kirkpatrick model. *J. Phys. A: Math. Gen.*, 27:5749–5772, 1994. Preprint cond-mat/9311016.
[9] Marinari E, Parisi G, and Ruiz Lorenzo JJ. Numerical simulations of spin glass systems. In [3], pages 59–98. Preprint cond-mat/9701016.
[10] Martin OC, Monasson R, and Zecchina R. Statistical mechanics methods and phase transitions in optimization problems. *Theor. Comp. Sci.*, 2001. Preprint cond-mat/0104428.
[11] Hertz J, Krogh A, and Palmer RG. *Introduction to the theory of neural computation.* Addison-Wesley (Redwood City, CA), 1991.
[12] Anderson PW, Arrow KJ, and Pines P, editors. *The economy as an evolving complex system.* Addison-Wesley (Reading, MA), 1988.
[13] Challet D and Zhang Y-C. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407–418, 1997. Preprint adap-org/9708006. See the URL http://www.unifr.ch/econophysics/minority/ for history and a large collection of commented references.
[14] Arthur WB. Inductive reasoning and bounded rationality (the El Farol problem). *Am. Econ. Rev.*, 84:406–411, 1994. Available at http://www.santafe.edu/~wba/Papers/Papers.html.
[15] Cavagna A, Garrahan JP, Giardina I, and Sherrington D. Thermal model for adaptive competition in a market. *Phys. Rev. Lett.*, 83:4429–4432, 1999. Preprint cond-mat/9903415.
[16] Challet D and Marsili M. Phase transition and symmetry breaking in the minority game. *Phys. Rev. E*, 60:R6271–R6274, 1999. Preprint cond-mat/9904071.
[17] Challet D, Marsili M, and Zecchina R. Statistical mechanics of systems with heterogeneous agents: Minority games. *Phys. Rev. Lett.*, 84:1824–1827, 2000. Preprint cond-mat/9904392.
[18] Marsili M, Challet D, and Zecchina R. Exact solution of a modified El Farol's bar problem: Efficiency and the role of market impact. *Physica A*, 280:522–553, 2000. Preprint cond-mat/990848.
[19] Garrahan JP, Moro E, and Sherrington D. Continuous time dynamics of the thermal minority game. *Phys. Rev. E*, 62:R9–R12, 2000. Preprint cond-mat/0004277.
[20] Challet D, Marsili M, and Zecchina R. Comment on "Thermal model for adaptive competition in a market". *Phys. Rev. Lett.*, 85:5008, 2000. Preprint cond-mat/0004308.
[21] Cavagna A, Garrahan JP, Giardina I, and Sherrington D. Reply to the comment by Challet, Marsili, and Zecchina. *Phys. Rev. Lett.*, 85:5009, 2000. Preprint cond-mat/0005134.
[22] Heimel JAF and Coolen ACC. Generating functional analysis of the dynamics of the batch minority game with random external information. *Phys. Rev. E*, 63:056121 (16 pages), 2001. Preprint cond-mat/0012045.
[23] Coolen ACC, Heimel JAF, and Sherrington D. Dynamics of the batch minority game with inhomogeneous decision noise. *Phys. Rev. E*, 2001. In print. Preprint cond-mat/0106635.
[24] Coolen ACC and Heimel JAF. Dynamical solution of the on-line minority game. Preprint cond-mat/0107600.
[25] Marsili M and Challet D. Continuum time limit and stationary states of the minority game. Preprint cond-mat/0102257.
[26] De Martino A and Marsili M. Replica symmetry breaking in the minority game. *J. Phys. A: Math. Gen.*, 34:2525–2537, 2001. Preprint cond-mat/0007397.
[27] Heimel JAF and De Martino A. Broken ergodicity and memory in the minority game. *J. Phys. A: Math. Gen.*, 34, 2001. In print. Preprint cond-mat/0108066.
[28] De Martino A and Marsili M. In preparation.
[29] Sherrington D and Kirkpatrick S. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35:1792–1796, 1975.
[30] Kosterlitz JM, Thouless DJ, and Jones RC. Spherical model of a spin-glass. *Phys. Rev. Lett.*, 36:1217–1220, 1976.

[31] Derrida B. Random energy model: an exactly solvable model of disordered systems. *Phys. Rev. B*, 24:2613–2626, 1981.

[32] Gross DJ and Mézard M. The simplest spin glass. *Nucl. Phys. B*, 240:431–452, 1984.

[33] Gardner E. Spin glasses with p-spin interactions. *Nucl. Phys. B*, 257:747–765, 1985.

[34] Crisanti A and Sommers H-J. The spherical p-spin interaction spin glass model: the statics. *Z. Phys. B*, 87:341–354, 1992.

[35] Kirkpatrick TR and Thirumalai D. p-spin interaction spin-glass models: connections with the structural glass problem. *Phys. Rev. B*, 36:5388–5397, 1987.

[36] Crisanti A, Horner H, and Sommers H-J. The spherical p-spin interaction spin glass model: the dynamics. *Z. Phys. B*, 92:257–271, 1993.

[37] Cugliandolo LF and Kurchan J. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Phys. Rev. Lett.*, 71:173–176, 1993. Preprint cond-mat/9303036.

[38] Franz S, Mézard M, Parisi G, and Peliti L. Measuring equilibrium properties in aging systems. *Phys. Rev. Lett.*, 81:1758–1761, 1998. Preprint cond-mat/9803108.

[39] Franz S, Mézard M, Parisi G, and Peliti L. The response of glassy systems to random perturbations: a bridge between equilibrium and off-equilibrium. *J. Stat. Phys.*, 97:459–488, 1999. Preprint cond-mat/9903370.

[40] Cavagna A. Irrelevance of memory in the minority game. *Phys. Rev. E*, 59:R3783–R3786, 1997. Preprint cond-mat/9812215.

[41] Challet D and Marsili M. Relevance of memory in minority games. *Phys. Rev. E*, 62:1862–1868, 2000. Preprint cond-mat/0004196.

[42] Sengupta AM and Mitra PP. Distribution of singular values for some random matrices. *Phys. Rev. E*, 60:3389–3392, 1999. Preprint cond-mat/9709283.

[43] Temesvári T, De Dominicis C, and Kondor I. Block diagonalizing ultrametric matrices. *J. Phys. A: Math. Gen.*, 27:7569–7595, 1994. Preprint cond-mat/9409050.

[44] Challet D, Marsili M, and Zecchina R. Modeling market mechanism with the minority game. *Physica A*, 276:284–315, 2000. Preprint cond-mat/9909265.

[45] Marsili M and Challet D. Trading behaviour and excess volatility in toy markets. *Adv. Compl. Sys.*, 4:3–17, 2001. Preprint cond-mat/0004376.

[46] Challet D, Chessa A, Marsili M, and Zhang Y-C. From minority games to real markets. *Quantitative Finance*, 1:168–176, 2001. Preprint cond-mat/0011042.

[47] Amit DJ, Gutfreund H, and Sompolinsky H. Statistical mechanics of neural networks near saturation. *Ann. Phys.*, 173:30–67, 1987.

[48] Müller B, Reinhardt J, and Strickland MT. *Neural networks: An introduction*. Springer-Verlag (Berlin), 1995.

[49] Weisbuch G, Kirman AP, and Herreiner DK. Market organization. In R. Conte, R. Hegselmann, and P.Terna, editors, *Simulating Social Phenomena*. Springer-Verlag (Berlin), 1997. See also the working paper "Market organisation and trading relationships" by the same authors, available at lps.ens.fr/˜weisbuch/papfeb.ps.