

NNT : 20XXIPPAXXXX

Thèse de doctorat



INSTITUT
POLYTECHNIQUE
DE PARIS



Navigating radical complexity

The influence of disorder, nonrelaxational dynamics
and learning on aggregate coordination

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à École polytechnique

École doctorale n°626 École doctorale de l'Institut Polytechnique de Paris (EDIPP)
Spécialité de doctorat : Physique

Thèse présentée et soutenue à Palaiseau, le 08/12/2023, par

JÉRÔME GARNIER-BRUN

Composition du Jury :

Leticia Cugliandolo Sorbonne Université	Rapporteur
Silvio Franz Université Paris-Saclay	Rapporteur
Marc Mézard Bocconi University	Examineur
Valentina Ros CNRS & Université Paris-Saclay	Examineur
Frédéric van Wijland Université Paris Cité	Examineur
Michael Benzaquen CNRS & École polytechnique	Directeur de thèse
Jean-Philippe Bouchaud CFM & Académie des Sciences	Co-directeur de thèse

Intentionally left blank.

Remerciements

J'aimerais commencer par remercier chaleureusement mes deux directeurs de thèse, Michael Benzaquen et Jean-Philippe Bouchaud. Dès le début, en acceptant de diriger ma thèse, et de manière continue tout au long de ces trois années, ils m'ont témoigné une très grande confiance qui a été essentielle pour moi. Leur disponibilité, patience et bienveillance m'ont accompagné depuis mon premier jour de stage, me permettant de m'épanouir scientifiquement et de ressentir un soutien constant tout au long de ma thèse. J'ai énormément appris à leur côté, et ce fut un réel privilège de pouvoir bénéficier de tels mentors.

Je remercie aussi les membres de mon jury de doctorat Leticia Cugliandolo, Silvio Franz, Marc Mézard, Valentina Ros et Frédéric van Wijland. Je suis particulièrement reconnaissant envers Leticia Cugliandolo et Silvio Franz d'avoir accepté d'endosser les rôles de rapporteurs. Je souhaiterais aussi remercier Frédéric van Wijland et Francesco Zamponi d'avoir constitué mon comité de suivi de thèse et pour les conseils qu'ils m'ont apportés.

My thesis was also enriched by a wonderful stay in Japan thanks first and foremost to Shin-ichi Sasa, who accepted to host me in his lab in Kyoto. I would also like to warmly thank Andreas Dechant, who immediately welcomed me and had the patience to teach me everything I know about stochastic thermodynamics. I also thank Kyogo Kawaguchi and Kyosuke Adachi for inviting me to Kobe on several occasions, as well as Takahiro Nemoto and Masato Itami for welcoming me and showing me around Kyodai.

Je souhaiterais ensuite remercier mes collègues et tous et toutes les membres de la chaire EconophysiX qui ont accompagné et souvent égayé ces trois années. Je remercie particulièrement Samy et Cécilia, évidemment, avec qui j'ai partagé la plus grande partie de cette aventure et sans qui de nombreuses périodes auraient assurément été beaucoup moins plaisantes, mais aussi Karl, Rudy, Salma et Nirbhay, pour les innombrables pauses à rallonge et discussions aussi variées qu'essentiels. J'aimerais également remercier chaleureusement mes "aïeux" José, Théo et Pierre, qui m'ont accueilli à mon arrivée et qui m'ont beaucoup appris, ainsi que les arrivants et arrivantes plus récents, en particulier Ruben et Antoine avec qui ce fut un plaisir de travailler. Je remercie aussi l'ensemble du personnel du LadHyX, qui s'est toujours montré extrêmement accueillant malgré les différences thématiques et ma présence intermittente sur le plateau de Saclay, et en particulier Daniel et Toaï pour leurs réponses à mes très nombreuses sollicitations.

Un très grand merci à mes colocataires, temporaires ou permanents, Théo, Thomas, Pierre, Mathilde et Valentine, qui ont partagé mon quotidien des trois dernières années et qui ont souvent su comprendre, sans pourtant pouvoir vraiment comprendre, mes nombreuses frustrations (ça y est, l'intégrale est à peu près résolue!). Je remercie aussi le reste de ma famille adoptive parisienne, Antoine,

Chloé, Coco, Raphaël et Romane, pour leur soutien mais peut-être surtout pour leur joie de vivre qui m’a souvent permis de souffler et de beaucoup rire. Je remercie évidemment tous et toutes les autres sudistes, ainsi que mes ami(e)s d’Angleterre et du master, qui ont aussi contribué à ce que cette thèse soit pour moi un chapitre heureux et épanouissant.

Je remercie enfin ma famille, sans qui je n’aurais pas pu faire ces longues études. Mes cousins, Florent, Sacha et Willy, qui sont pour moi comme des frères, et qui se sont toujours montrés à l’écoute. Mes sœurs, Aurélie et Julia, bien-sûr, qui sont les premières à me soutenir et à m’encourager, et pour qui j’ai une énorme admiration qui m’a toujours poussé à essayer d’être la meilleure version de moi-même. Mes parents, enfin, qui m’ont transmis la curiosité et m’ont sans cesse écouté et soutenu depuis mon premier jour. Merci pour tout.

Je dédie ce manuscrit à ma grand-mère Andrée qui, bien qu’elle ait dû arrêter l’école à contrecœur après son certificat d’étude, a su me faire comprendre très jeune la chance que représentait les études et ce qu’était le plaisir d’apprendre. Elle était pour moi un exemple d’humilité et d’ouverture d’esprit.

Foreword

The journey behind this thesis starts on a cold and rainy day of December 2017 in London. Then a student at Imperial College London predestined to a career in engineering, I met with Michael Benzaquen to discuss an internship opening in his research group. As I was looking for a research experience in hydrodynamics, I had stumbled on his webpage a few weeks earlier. There, internship and PhD proposals broadly combining physics and economics had piqued my interest. Due to my background in fluid mechanics, I ended up studying microfluidics with Gabriel Amselem and himself at LadHyX. It would involve two months of painstaking experiments, creating and deforming hundreds of droplets of around half a millimeter in diameter in fragile silicon canals, but I had a foot in the door.

Towards the end of the internship, Gabriel – to whom I had mentioned my interest for interdisciplinary science – told me about a seminar that was to be given by Jean-Philippe Bouchaud, titled “De la Physique Statistique aux Sciences Sociales”. From there, I left with not only a clearer picture of what “Econophysics” was, but also with the conviction that, indeed, the sudden switch from a prosperous economy to a recession is not so different to liquid water turning into an ice cube. In any case, my mind was set, I needed to become a statistical physicist to eventually study these complex systems.

Fast forward to a couple of years later. In order to find a good internship and complete my “Physics of Complex Systems” master in Paris, I came back to Michael, initially just to ask about the different statistical physics labs in the Paris area. Of course, the idea of joining the then new *EconophysiX* chair had crossed my mind, so when Michael mentioned a project co-supervised by Jean-Philippe and linking portfolio optimization to the theory of spin-glasses, I did not hesitate much – nonetheless making sure no microfluidic experiments were involved. Although I did not initially know much about spin-glasses, their incredibly rich phenomenology hidden behind a deceptive simplicity quickly caught my interest, and the internship confirmed my motivation to pursue a PhD. Yet, before embarking in these further studies, I had planned to explore another aspect of statistical mechanics (and another region of the world) through a research experience in Japan. Following the advice of my then teacher Prof. van Wijland, I had my eyes set on Prof. Sasa’s lab in Kyoto to work on stochastic thermodynamics. This was March 2020, however, so the project was cut short for a rather evident reason. So be it, I had found a great lab and a research topic that I very much liked, and it was with enthusiasm and excitement that I started my PhD in October 2020.

During the past three years, I was fortunate enough to work on a variety of projects – some successful, others less so –, touching upon different aspects of both socioeconomic modeling and of statistical mechanics. Clearly, a large part of the thesis is rooted in the theory of disordered systems, in the continuity of

my very first work associating spin-glasses and optimal portfolios. Leveraging and extending results from spin-glasses towards the specificities of socioeconomic systems is, in my opinion, the main contribution of the thesis. This being said, other problems outside of the realm of disordered systems encountered along the way have proved to be extremely useful for my understanding of the challenges of socioeconomic modeling and of the possible answers (or lack of) provided by statistical mechanics.

Academic research being a collective phenomenon at heart, my PhD involved many interactions. First and foremost with my immediate colleagues of course, but also with many inspiring fellow students and scientists at schools in Cargèse, Beg Rohu or les Houches. Towards the end of my thesis, I also had the chance to work hand in hand with others on two collaborative projects. First in Prof. Sasa's group in Kyoto, where I eventually had the amazing opportunity of going for a three month period in 2022. There, I benefited from countless discussions with Andreas Dechant, and of his and Prof. Sasa's expertise, providing me with the ideal introduction to stochastic thermodynamics. Then, in my final year within the *EconophysiX* research group with Ruben Zakine and Antoine-Cyrus Becharat, allowing me to learn from Ruben's competence in active matter and to guide, to the best of my ability, Antoine-Cyrus in the first year of his PhD. The outcome of these collaborations is partially included in the manuscript, with my coworkers' agreement of course.

In addition to the diverse and stimulating scientific projects mentioned above, my PhD also encompassed other responsibilities integral to the learning experience of becoming a researcher. As a teaching assistant, on the one hand, at ENSAE and École polytechnique and for courses ranging from the first to final year of studies. Within the *EconophysiX* research group on the other, as I was tasked with organizing and chairing the monthly group meetings from the first year onwards. Both these assignments turned out to be quite enjoyable and, in my opinion, beneficial to communicate my research effectively and expand my scientific culture.

As hinted above, the present manuscript is predominantly written as a statistical physics work. Indeed, all of my publications have been in physics journals, and our toy-modeling approach is undoubtedly different from what would be found in typical economics research. Nonetheless, I would like the ideas exposed here to be understandable to non-physicists. As a result, without being self-contained, a relatively significant part is dedicated to theoretical background, which might perhaps appear unnecessary to the initiated reader. If so, the dedicated chapter (Chap. 2) can of course be skipped, or only consulted when required (explicit references will be made to the relevant subsections).

List of publications

- [1] **G.-B. J.**, Benzaquen M., Ciliberti S., & Bouchaud J.-P. (2021). “A new spin on optimal portfolios and ecological equilibria”. *J. Stat. Mech.* 093408 (featured in Highlights)
- [2] **G.-B. J.**, Bouchaud J.-P. & Benzaquen M. (2022). “Bounded rationality and animal spirits: a fluctuation-response approach to Slutsky matrices”. *J. Phys. Complex.* 4, 015004 (special issue *Focus on Fundamental Theory and Application of Complexity Economics*)
- [3] **G.-B. J.**, Benzaquen M. & Bouchaud J.-P. (2023). “Unlearnable Games and Radical Complexity: A Simple Model for a Complex World”. *In prep.*
- [4] Dechant A., **G.-B. J.** & Sasa S.-i. (2023). “Thermodynamic bounds on correlation times”. *Phys. Rev. Lett.* 131, 167101 (featured in Editor’s Suggestion)
- [5] Zakine R., **G.-B. J.**, Becharat A.-C. & Benzaquen M. (2023). “Socioeconomic agents as active matter in nonequilibrium Sakoda-Schelling models”. *Submitted to Phys. Rev. E*, arXiv:2307.14270
- [6] Alemany Juvanteny I., Rose J., **G.-B. J.**, Scott A. & Doorly D. (2022). “Random walk diffusion simulations in semi-permeable layered media with varying diffusivity”. *Sci. Rep.* 12(1), 10759

Contents

I	Motivation and background	1
1	Introduction	3
1.1	From spin-glasses to complex systems	5
1.1.1	A bit of history	5
1.1.2	Main features	6
1.1.3	The spin-glass <i>cornucopia</i>	7
1.2	Agent-based modeling in socioeconomics	8
1.2.1	Agents as atoms	9
1.2.2	Utility theory	9
1.2.3	Homo economicus, bounded rationality and the logit rule .	10
1.2.4	Simplicity versus fidelity	13
1.3	Toy stories	14
1.3.1	The Schelling model, micromotives and macrobehaviours .	14
1.3.2	The Random-Field Ising Model	17
1.3.3	A word on universality	20
1.4	Thesis layout	22
2	Theoretical foundations	25
2.1	Equilibrium statistical mechanics	25
2.1.1	Thermodynamics and the Gibbs-Boltzmann measure	25
2.1.2	Detailed balance	29
2.1.3	Langevin dynamics	30
2.1.4	A word on the relaxation time	32
2.1.5	Ergodicity	32
2.2	Disordered systems	33
2.2.1	The replica method	34
2.2.2	TAP approach	36
2.2.3	Complexity	37
2.2.4	Replica symmetry breaking	39
2.2.5	Dynamics	40

Contents

2.3	Summary of the key notions	42
II Radical complexity and detailed balance violation in socioeconomics		45
3	Rationality versus complexity: the example of portfolio optimization	47
3.1	The portfolio optimization problem	48
3.1.1	Problem statement	49
3.1.2	Link with population dynamics	50
3.1.3	Spin-glass formulation	51
3.1.4	Empirical data	53
3.2	Numerical experiments	54
3.2.1	Exact enumeration	54
3.2.2	“Pruning” algorithm	56
3.3	Analytical setup	59
3.3.1	Self-consistent equation	59
3.3.2	Complexity and number of solutions	61
3.4	Distribution-specific results	63
3.4.1	Gaussian disorder	63
3.4.2	Uniform disorder	65
3.4.3	Bridging the gap: generalized normal distribution	67
3.5	Disorder chaos	70
3.6	Conclusion	72
4	Slutsky matrices and the necessity of a global utility	77
4.1	Consumer choice theory	77
4.2	Thermodynamics of bounded rationality	79
4.2.1	A single agent	79
4.2.2	A fluctuation-response relation	80
4.2.3	Many agents	81
4.2.4	Aggregate Slutsky matrices	83
4.2.5	Near-rational limit $\beta \rightarrow \infty$	84
4.3	Animal spirits	86
4.3.1	Interactions and herding	86
4.3.2	Concentration for $\beta \rightarrow \infty$	88
4.3.3	Finite β effects	90
4.4	Consequences on the Slutsky matrix	90
4.5	Discussion	95
4.5.1	Global vs. individual utilities	95

4.5.2	Equivalence of ensembles	97
4.6	Conclusion	99
III	A unifying disordered model: the SK-game	103
5	Presentation of the model and numerical results	105
5.1	Motivation	105
5.2	A simple model for a complex world	107
5.2.1	Set-up of the model	107
5.2.2	The interaction matrix	109
5.3	Overview and numerical results	111
5.3.1	A word on the average reward	111
5.3.2	Phase diagram in the noiseless limit	111
5.3.3	Noisy learning	114
5.3.4	Distribution of individual rewards	116
5.3.5	Unpredictability of equilibria	119
5.3.6	Increasing cooperativity	120
5.3.7	Self-reinforcement and habit formation	120
5.4	Core message	121
6	Detailed analysis of the model	125
6.1	Fixed point analysis and complexity	125
6.1.1	Critical noise level	126
6.1.2	The elusive complexity	127
6.2	Counting limit cycles	129
6.2.1	Cycles without memory	129
6.2.2	Cycles with memory	130
6.3	Dynamical Mean-Field Theory	131
6.3.1	Derivation	133
6.3.2	Numerical integration	134
6.3.3	Interpretation in a socioeconomic context	135
6.4	Noiseless learning	136
6.4.1	The memory-less limit	136
6.4.2	Memory helps convergence to fixed points	137
6.4.3	Anomalous stretching of cycles	139
6.5	Noisy learning	142
6.5.1	(Quasi-)fixed points	143
6.5.2	Memory onset transition	145
6.5.3	Aging	146
6.5.4	Chaos and (quasi-)limit cycles	148

Contents

6.6	Conclusion	151
6.6.1	Blindsided by complexity	151
6.6.2	Technical results and conjectures	152
6.6.3	Extensions and final remarks	153
IV	Detailed balance violation in the absence of disorder	157
7	A Sakoda-Schelling model with no neighborhoods	159
7.1	Motivation	160
7.2	Presentation of the model	161
7.3	In or out of equilibrium?	162
7.4	Numerical experiments	163
7.4.1	Phase separation	164
7.4.2	Critical point and exponents	165
7.5	Mean-field description	167
7.5.1	Linear stability analysis	168
7.5.2	Comparison with numerical simulations	169
7.5.3	Local move approximation	169
7.5.4	Alternative derivation and fluctuating hydrodynamics	171
7.6	Generalized thermodynamic mapping	173
7.7	Two population extension	175
7.8	Conclusion	177
8	Steady-state preserving out-of-equilibrium currents	181
8.1	General idea	182
8.1.1	Physical setup	182
8.1.2	Relaxation to the NESS	183
8.1.3	Barrier crossing out-of-equilibrium	184
8.2	Theoretical bounds on the correlation time	185
8.2.1	General variational statement	185
8.2.2	Dissipation speed limit	189
8.2.3	Geometric speed limit	190
8.3	Accelerating relaxation in practice	191
8.3.1	Ichiki-Ohzeki forcing	191
8.3.2	An extension to complicated currents	193
8.3.3	Perspectives	196
8.4	Conclusion	197

V Conclusion and future works	201
References	207
Appendices	235
A Optimal portfolios	235
A.1 Full self-consistent equation	235
A.2 Sommerfeld-like expansions	236
A.3 Reaction term	237
A.4 Detailed resolution of the characteristic equations	238
A.5 Generalized normal distribution	242
A.6 Heterogeneous returns and growth rates	243
B Slutsky matrices	247
B.1 General “thermodynamic” relations	247
B.1.1 Single agent Slutsky matrix	247
B.1.2 Many agent Slutsky matrix	248
B.1.3 Fluctuation-response relations	249
B.1.4 Aggregate Slutsky matrix	250
B.1.5 Gaussian fluctuations	251
B.2 Numerical methods	253
B.2.1 General idea	253
B.2.2 Metropolis-Hastings acceptance rate	253
B.2.3 Computing the Slutsky matrix	254
B.3 Interacting model	254
B.3.1 Non-interacting limit – canonical ensemble	254
B.3.2 Finite interactions – grand-canonical ensemble	256
B.3.3 Slutsky matrix in the $\beta \rightarrow \infty$ limit	259
B.3.4 Equivalence of ensembles	262
B.4 A Hamiltonian utility function	265
C SK-game	267
C.1 Static NMFE	267
C.2 Fixed point complexity with cooperativity	268
C.3 Limit cycle complexity with memory	271
C.3.1 Fixed points	278
C.3.2 Two-cycles	278
C.4 Derivation of the DMFT equations	280

Contents

D Schelling model	283
D.1 Lyapunov function for non-local moves	283
D.2 Local mean-field description and LSA	284
D.3 Linear stability for two coupled populations	285
D.4 LSA for two populations with local moves	287

Part I

Motivation and background

Chapter 1

Introduction

The difficulty lies not so much in developing new ideas as in escaping from old ones.

John Maynard Keynes

The idea of drawing parallels between collective human behavior and physics is not new. Starting from Daniel Bernoulli¹ [7], who is most famous for his contributions to hydrodynamics, theoretical economics has indeed long attempted to establish laws, akin to the laws of physics, to describe markets, consumption or production.

In this endeavor, the pioneers of neoclassical economics such as Léon Walras initially found inspiration in mechanics and the rigorous mathematical formalism around it [8]. Whether realistic or not, by considering individuals as egoistical, non-interacting and clairvoyant, human behavior can indeed be put into deterministic equilibrium equations where supply perfectly matches demand. From then on, economic theory has mostly evolved in a very axiomatic fashion, with a strong preference towards very simple models that allow for rigorously proved theorems.

Interestingly, it is then theoretical physics that ventured back into socioeconomics following the major advances of statistical mechanics in the XXth century. The discovery of *universality*, whereby the macroscopic properties of large classes of systems turn out to be independent of their microscopic details, and the related effectiveness of simplified *toy models* to describe complex phenomena have indeed made the application of newly developed methods to socioeconomic problems very compelling. Collective human systems such as the economy being comprised of a large number of interacting individuals and prone to display abrupt changes such as crises, these tools naturally appear quite adequate.

¹Not to be mistaken with his uncle Jacob Bernoulli, known for the Bernoulli numbers and distribution, or his father Johannes specialized in infinitesimal calculus or any of the many other Bernoulli mathematicians.

Chapter 1. Introduction

Being entirely at odds with the *representative agent* paradigm that has been at the center of many classical theories, the fields of “Econophysics” and “Sociophysics” have encountered mixed success overall. On one hand, a very large number of interesting models have emerged (see e.g. [9–15]), some of which replicate essential empirical observations such as power law distributions [16]. In finance, in particular, where data is plentiful and theories can be readily tested, some tools from statistical physics have had a significant impact [17, 18]. On the other hand, even some of the most elementary concepts in statistical mechanics have not well permeated to mainstream economics, which is still extremely dependent on general equilibrium models that rely on the typical assumptions of a rational representative agent. This being said, the failure of macroeconomics to predict the 2008 crisis has somewhat accelerated the acceptance that heterogeneity and feedback loops might be key to understand and predict large scale complex phenomena [19–22]. There is therefore still some hope for socioeconomics to embrace what theoretical physics has to offer [23, 24].

The opportunity for ideas from statistical physics to finally make a breakthrough in socioeconomics is also driven by the success of the theory of disordered systems in other fields. Indeed, the theory of spin-glasses, initiated to study the effect of defects on magnetic materials roughly 50 years ago, has bloomed, through its interaction with other disciplines, into complex systems science. Culminating with Giorgio Parisi’s Nobel Prize in 2021, spin-glasses have been an extremely fruitful playground both in terms of theory and applications, being simultaneously very general and simple in their statement, and incredibly profound and far-reaching in their consequences. In a nutshell, the theory of disordered systems has demonstrated that heterogeneous interactions between a large number of constituents have highly non-trivial effects, both on the static and dynamic properties of the system. *A priori*, many socioeconomic models should fall within this description.

The overarching goal of this thesis is to use and expand results from disordered systems towards socioeconomic systems, which will also lead us to study some specific problems using equilibrium and out-of-equilibrium statistical mechanics along the way. In this introduction chapter, we will first present some very basic phenomenology from spin-glasses and complex systems science. Following this rather pedestrian overview, we will lay the ground rules on the socioeconomic side to give an idea of the typical modeling assumptions that will be made in this thesis. Finally, two prototypical examples of existing socioeconomic “toy models” will be introduced to illustrate first what we mean by a toy model in this context and second the two-way interaction that can exist in the modeling of human behavior with physics-inspired approaches.

1.1 From spin-glasses to complex systems

1.1.1 A bit of history

Despite its apparent disconnection with physics, much of what will be presented in this thesis can be traced back to good old fashioned laboratory work. In the 1970s, an intriguing puzzle emerged from the study of Cu-Mn and Au-Fe alloys. In these alloys, where a small fraction of moment-bearing elements (here Mn and Fe) is incorporated in a non-magnetic host (Cu and Au), a sharp but non-divergent peak in magnetic susceptibility can be identified for a concentration-dependent critical temperature [25, 26]. Surprisingly, the magnetization of the material is not non-zero below the critical temperature as in a standard magnetic sample, in which the susceptibility diverges at the transition. Instead, one can observe the emergence of order, in the sense that the local polarization of magnetic atoms stays correlated in time, but the orientation of individual atoms still appears to be random. Understanding how these “spin-glasses”² can display such frozen yet disorganized low-temperature states then became a central theoretical question.

The answer to this mystery can be explained in relatively simple terms. In the Cu-Mn or Au-Fe alloys, the interaction between polarized atoms oscillates in sign with the distance between the particles. As a result, when there is a relatively small fraction of the moment-bearing particles that is *randomly distributed* in the non-magnetic host matrix, there will be no mean ferromagnetic or anti-ferromagnetic (tendency for spins to align or anti-align respectively) interaction. In the low temperature region, each spin will therefore align in the direction that is favored by its local random environment. Importantly, the heterogeneity of the sample gives rise to a phenomenon known as *frustration*, illustrated in Fig. 1.1. Indeed, most sites will have contradicting interests with each other, explaining why no single direction can emerge from the oscillating interactions. Note that the effectively random nature of the interactions is why spin-glasses are classified as *disordered systems*, which are typically understood to include a wider range of models that are not necessarily describing spin-like quantities.

To describe this situation analytically and as generically as possible, Edwards and Anderson proposed a simple Hamiltonian to describe the energy of spin-glasses in 1975 [27]. In their model, the energy of a given configuration is simply the Ising Hamiltonian

$$\mathcal{H} = - \sum_{(i,j)} J_{ij} S_i S_j, \quad (1.1)$$

where, importantly, $J_{ij} = J_{ji}$ the interaction between the spins on sites i and j

²Here, the term glass refers to the conceptual similarity between the seemingly random polarization of atoms and the positional disorder characteristic of conventional glasses.




Figure 1.1: Illustration of the frustration phenomenon that arises in disordered systems. Solid links indicate $J_{ij} = J_{ji} = 1$ while dashed links represent $J_{ij} = J_{ji} = -1$. Left: non-frustrated “ferromagnet”, all spins can be happily aligned. Right: frustrated configuration, spin C is undecided as it cannot simultaneously align with B and anti-align with A .

is taken to be *random* (with usually a zero mean), while the S_i , $i = 1, \dots, N$, are standard binary spins, $S_i = \pm 1$.

In their original paper, Edwards and Anderson (EA) considered nearest neighbor interactions, and showed the existence of a spin-glass phase below the critical temperature. As will be detailed later, the fully connected version of the model due to Sherrington and Kirkpatrick (SK) [28] is simpler – or rather not as difficult – to study analytically and becomes particularly relevant for interdisciplinary applications.

1.1.2 Main features

The theoretical understanding of these idealized spin-glasses then made tremendous progress in the decades following their introduction. While some of these results will be exposed in more detail in the next chapter, let us summarize the main features that emerge from the toy models introduced by EA and SK.

First, these simplified models present very singular static properties as a result of the frustration illustrated in Fig. 1.1:

- A very large number of locally optimal spin configurations,
- These configurations are near-degenerate (multiplicity of quasi-equivalent solutions) and separated by enormous energy barriers,
- The local optima are extremely sensitive to the interaction couplings, i.e. solutions are “fragile”.

These unusual static properties then have important dynamic consequences, namely:

- A very strong dependence on initial conditions,
- Extremely slow (sometimes infinitely slow) relaxation dynamics and ergodicity breaking.

The fact that such simply stated models can have this extremely rich and complex phenomenology is remarkable in itself. One could even argue that the theoretical models of spin-glasses are in some sense more interesting than the physical systems they originally aimed to describe, as they require an entire new way of thinking about thermodynamics (due to their non-ergodic nature). Whether this entire phenomenology is realistic in the context of metallic alloys is then of secondary importance. As a matter of fact, it turns out that while both the EA and SK model recover a finite susceptibility peak at the critical temperature, they also give rise to a cusp in the specific heat, which has not been identified in experiments.

1.1.3 The spin-glass *cornucopia*

It should now be clear that the merit of theoretical models of spin-glasses does not lie in their ability to accurately describe real physical systems. If this is the case and that they are merely an intellectual curiosity, then what is the fuss all about, and more importantly why would the Swedish Academy of Sciences award Giorgio Parisi a Nobel prize? (And what's more shared with two climate scientists who's work then appears much more useful and relevant.)

The answer lies in what Philip Anderson elegantly called the spin-glass *cornucopia*³ [29]. While being a relatively poor model of Cu-Mn alloys and of anything too precise, spin-glasses are a great model for many other things.

First and foremost, the frustration characterizing spin-glasses means that finding the ground state, i.e. the spin configuration minimizing a given Hamiltonian, is what is known in optimization as an “NP-hard” problem. In other words, there is no known algorithm that systematically finds the configuration $\{S_i = \pm 1\}$ yielding the lowest possible value of $\mathcal{H}$ defined in Eq. (1.1) in less than $\sim 2^N$ steps. While the statistical mechanics of this spin-glass will not give the solution for a specific instance of the disorder, it can provide precious information on the *typical* outcome one may expect. One may then use these insights to understand some key aspects of other “NP-hard” optimization problems, such as graph partitioning [30] or the famous traveling salesman problem [31,32]. The most notable contribution from theoretical physics to these optimization problem is perhaps the heuristic algorithm known as simulated annealing, introduced by Scott Kirkpatrick (of the SK model) and coworkers in 1983 [33]. By slowly cooling the system from above the critical temperature, the lower energy states can be reached with much higher probability than if starting directly in the very low temperature region. While this is not true for all optimization problems, the technique can be very effective in some cases.

³“Horn of plenty” in latin, i.e. a plentiful supply in this context.

Chapter 1. Introduction

Beyond these rather natural links towards optimization – the Hamiltonian being only a specific instance of objective function to minimize – other fruitful interdisciplinary extensions of spin-glasses have also emerged. Most notably, the theory of both biological and artificial neural networks has been, and still is, an extremely successful playing field for spin-glass physicists. Starting from the seminal work of Little [34] and Hopfield [35], and subsequently of Gardner and Derrida [36], the modeling of neuron activations as binary variables has given a very natural application for the theory of spin-glasses. Given the poor theoretical understanding of deep neural networks despite their popularity and effectiveness, it appears likely that there is still a lot of potential for the application of spin-glass theory to neural networks. The recent success of “transformers”, which can be understood with modern Hopfield networks [37], and of diffusion models, for which the physical interpretation is rather natural [38], also give very promising perspectives in this direction. Other applications where spin-glasses have proved to be conceptually important include evolutionary biology [39] and population dynamics [40, 41], as well as optimal portfolios [42] and agent-based models [43] in the spirit of this thesis.

In the spin-glass community, these interdisciplinary applications are sometimes referred to as the “beyond”, in reference to Mézard, Parisi and Virasoro’s now classic book “Spin glass theory and beyond” [44]. Amusingly, while this “beyond” was originally restricted to the optimization and biological problems mentioned above, the very first page of the book introduces the concept of frustration illustrated in Fig. 1.1 with the relationship between three individuals. Clearly, it is therefore natural to imagine that frustration and heterogeneity play a role in the collective behavior of socioeconomic systems. As a matter of fact, a chapter dedicated to socioeconomic applications (and mentioning some of the results presented in this thesis) is now included in the recently edited sequel, “Spin glass theory and *far* beyond” [45], which compiles the contribution of over a hundred authors.

1.2 Agent-based modeling in socioeconomics

Before making the link between statistical physics and any socioeconomic system, some clarification is required as to what exactly is meant by the latter. Indeed, both economics and sociology encompass a very wide range of approaches, some of which are not necessarily quantifiable as such. Whether one is thinking of political economics or the psychological (or even physiological) theories at the root of human behavior, the possibility of drawing parallels with theoretical physics is not evident at all. Let us define more precisely the prerequisites for a socioeconomic problem to be somewhat within the realm of the discipline in which this thesis lies.

1.2.1 Agents as atoms

First and foremost, it is important to highlight that attempts at using statistical mechanics very often imply the presence of a large number of constituents in the system. Instead of atoms or colloids, the elementary entity in this context will often (but not always) be an *agent*. In practice, an agent can be very different things: an individual person or household obviously, but also a firm, an organization, a central banks etc.

A popular idea in classical economics is then that of the *representative agent*. To avoid going through what we will see is the painstaking process of accounting for heterogeneity among a population for instance, the idea is to describe a model or situation with a very small number of agents (often just a single “representative consumer” and a handful of other actors), acting as proxies for the entire economy. While this approach is not necessarily inherently wrong⁴, a large part of – if not all – potential emergent behavior can very easily be lost. As indirectly admitted by Nobel prize laureate Robert Lucas in a somewhat surreal attempt at defending the failures of macroeconomics [46], a consequence of this paradigm is that the 2008 economic crisis was not predicted because economic theory predicts that such events cannot be predicted. Indeed, by removing collective effects and feedback loops, only exogenous shocks can typically destabilize a system under this representative assumption, clearly contradicting empirical facts, see e.g. [47, 48] for a comprehensive discussion on the topic.

To move away from this oftentimes simplistic reasoning, there has been significant efforts since 2008 to develop genuine agent-based models, with a large number of heterogeneous entities interacting to a certain degree. It is on these type of models that we will mostly focus in this work and where statistical physics likely has the largest part to play.

1.2.2 Utility theory

Regardless of the heterogeneous or representative nature of the agent considered, a key and unavoidable step in any modeling endeavor then requires putting a plain word description in equations or systematic rules. Reasoning in terms of first principles, as one would do in physics, is not really possible when considering socioeconomic systems as there are no evident universal or elementary rules (akin for example to energy minimization), and is further complicated by the difficulty of conducting repeatable behavioral experiments on humans and organizations.

As a major simplification, utility theory consists in postulating the existence of a so-called utility function $u(X)$, where here X would be (in a loose sense) a state of the world, part of the *set of possible alternatives* offered to the agent [49].

⁴We will see how this idea can be revisited with Dynamical Mean-Field Theory in Chap. 6.

Chapter 1. Introduction

In a typical consumer decision problem, the set of X are the possible consumption choices for example (e.g. banana, apple and orange). The utility function essentially quantifies⁵ the agent's satisfaction following a given decision. Usually the state X is restricted to things that might vary following the agent's action, and can then be viewed as an agent-specific analogue of the energy. Whereas energy is minimized, it is assumed that the agent maximizes his or her utility.⁶ Importantly, the central assumption behind the utility is that an agent will prefer the option with the highest utility if given a choice between two outcomes. To go back to the previous subsection, one of the major differences between the representative agent paradigm and the agent-based models of interest here is that the utility in the latter is strongly affected by the choice of others, and not only by the choices of oneself as in the former.

Within this framework, economists then usually make further postulates on the mathematical properties of the function u [51]. Not without consequences, when X is some measure of consumption, $u(X)$ is taken to be strictly increasing for example: more is always better and there is no satiation [49]. This hypothesis is of course extremely debatable, particularly in today's context where excessive consumption is under scrutiny for obvious reasons, and should therefore be kept in mind before making rash conclusions. To account for the fact that satisfaction is (should be?) somewhat diminishing as goods are accumulated, the utility function is also often assumed to be concave, which is known as the law of diminishing marginal utility. In any case, the utility function will be explicitly given in the specific problems we will consider, clearly stating the underlying assumptions that are made.

1.2.3 Homo economicus, bounded rationality and the logit rule

Provided one accepts utility theory as the best way to formalize some measure of satisfaction for agents, there remains the question of how agents use such a hypothetical quantity to make their decision. Given the assumption that an individual should always go for the outcome leading to the highest value of utility given the choice, the most direct approach is to take agents to be strict optimizers, determining their next step as the one that will maximize their utility gain.

This “homo economicus” paradigm is the most common in conventional economics, despite some contradicting empirical evidence. Beyond any behavioural

⁵To be perfectly precise, we are speaking here of a *cardinal* utility function, in contrast with so-called *ordinal* utility functions that only rely on ordered preferences without having to precisely quantify levels of utility.

⁶Note that the mirroring between utility and energy is likely not entirely coincidental: the origins of utility theory can be traced back to the physicist Daniel Bernoulli [7], while Irving Fisher, who significantly contributed to its development, was trained by Willard Gibbs [50].

experiment one could undertake to prove or disprove that human beings systematically improve their gain when possible, we will shortly see that this scenario can also be challenged from a computational point of view. If there are many near-degenerate solutions to the optimization problem at hand (like in a spin-glass), then it appears unlikely for anyone to immediately find the *optimum optimum* in order to make their next decision – as it might be an “NP-hard” task! In many cases, evidence points rather towards the unconscious use of powerful heuristics [52], although the precise decision making process going on in one’s mind is of course difficult, if not impossible, to decipher. In any case, it appears relevant to introduce some variation to the strict optimizer belief.

There are several ways to relax such an assumption. One is that agents have a limited attention and cannot process all the information accessible to them, see e.g. [53] and refs. therein. As argued by Gabaix, this may effectively lead to *perceived* prices that differ from real prices for example. Another traditional line of thought in the literature on choice theory is to replace the deterministic utility optimization prescription by a *stochastic* choice rule: the utility will only be maximized with a certain probability [54]. In some sense, this is a very general approach, as the stochasticity in the decision making could be interpreted to include a wide range of possible mechanisms.

A common model for this probabilistic decision making is the so-called “logit” rule [54, 55]. Given a set of choices $\mathcal{A}$, it states that the probability for an agent to pick the option $\alpha \in \mathcal{A}$ is given by

$$P_\alpha = \frac{e^{\beta u_\alpha}}{\sum_{\gamma \in \mathcal{A}} e^{\beta u_\gamma}}, \quad (1.2)$$

where β is a parameter known as the “intensity of choice”, or more loosely as the rationality of the agent, and u_α is simply the utility associated to choice α . In the limit $\beta \rightarrow \infty$, one then recovers the “homo economicus” picture, as

$$\lim_{\beta \rightarrow \infty} P_\alpha = \begin{cases} 1 & \text{if } u_\alpha = \max_{\gamma \in \mathcal{A}} u_\gamma, \\ 0 & \text{otherwise.} \end{cases} \quad (1.3)$$

On the other hand, the limit $\beta \rightarrow 0$ will yield a completely random behavior in which all options are equiprobable. As summarized in Ref. [15], there are several ways to justify this proposition from classical choice theory.

The first is to imagine that the probabilistic nature of the rule stems from the difference between the true utility of a choice u_α and the perception of the agent making the decision, which we will write $\hat{u}_\alpha$. Taking the so-called “random utility model” [56]

$$\hat{u}_\alpha = u_\alpha + \epsilon_\alpha, \quad (1.4)$$

Chapter 1. Introduction

the logit decision rule is recovered if the ϵ_α are i.i.d. Gumbel random variables [57]. A natural question is then, why would these perturbations ϵ_α follow a distribution which is more commonly encountered when dealing with extreme value statistics than behavioral sciences? The short (but unsatisfying) answer to this question is that it gives the logit rule, or in other words that this perceived utility idea is not a great justification. A more complicated but interesting explanation is given in a modest footnote of [58] (pp 32-33). In a nutshell, if the agent strictly maximizes a $\hat{u}_\alpha$ that is now affected by new variables that are themselves Gaussian distributed, then the extreme value statistics resulting from the maximization over these other “hidden” random parameters allows one to recover the Gumbell distributed ϵ_α . Despite this more subtle argument, the justification remains quite weak.

The second approach is to introduce an information cost for the agents, which is taken to be a decreasing function of the entropy

$$S = - \sum_{\gamma \in \mathcal{A}} P_\gamma \log P_\gamma. \quad (1.5)$$

Indeed, the entropy is a rather natural measure of the precision of a distribution, as it is maximized for uniform P_γ and minimized for an exact choice $P_\gamma = \delta_{\alpha,\gamma}$. Maximizing the total expected gain with this cost included

$$\mathcal{L}[P] = \sum_{\gamma \in \mathcal{A}} P_\gamma u_\gamma - \mathcal{C} \left(- \sum_{\gamma \in \mathcal{A}} P_\gamma \log P_\gamma \right), \quad \text{with} \quad \sum_{\gamma \in \mathcal{A}} P_\gamma = 1, \quad (1.6)$$

one then recovers the logit rule, where $\beta^{-1} = -\mathcal{C}'(S) > 0$ can now be interpreted as the *marginal utility cost of information* [59]. Within this interpretation, the decision-making process is essentially an exploration-exploitation compromise [60].

A final way to justify the decision rule is to take the axiomatic road. Indeed, Eq. (5.1) satisfies the axiom of *independence from irrelevant alternatives* (part of the larger *choice* axiom [54]), meaning that the ratio P_α/P_γ , $\alpha \neq \gamma$, is unaffected by the introduction of a supplementary choice. Interestingly, having Gumbel distributed ϵ_α in the random utility model is the only way to satisfy this axiom, giving a bit more context to the first justification [58].

Naturally, being able to confirm or infirm either the logit rule itself or the underlying axioms in experimental studies would be most desirable. Overall, most empirical attempts have been either inconclusive, or ruled somewhat against the traditional axioms of choice theory [61], see [62] for a more recent discussion. This being said, it can be difficult to decipher whether classical choice theory is to blame, or if utility theory itself is an underlying limitation.

Despite this discouraging experimental picture, it is important to keep in mind that most economic models do not even consider bounded rationality, and still rely

on strict utility maximization. While the logit rule may not be a perfect (or even good) model for decision making, it is already a conceptually very important development from what has traditionally been done. In that respect, the most important justification for its use may just be the fact that it will allow us to carry out many computations, as will become clear in the next chapter (and is likely already clear for the reader familiar with statistical mechanics). As nicely summarized by its father R. Duncan Luce [61], “*Despite [...] empirical difficulties, there remains a tendency to invoke the choice axiom [which leads to the logit rule] in many behavioural models – often implicitly. This is partly because it is so simple and the resulting computations are so easy. Perhaps the greatest strength of the choice axiom, and one reason it continues to be used, is as a canon of probabilistic rationality.*”

Let us mention a final interesting approach to bounded rationality, which will be relevant at later stages of this thesis. Shortly after pioneering the concept of bounded rationality [63], Herbert Simon proposed the idea of *satisficing* solutions [64]. The word, which is a contraction of *satisfying* and *sufficing*, refers to the fact that most people are likely to be content with a choice that matches or exceeds their expectation, without necessarily being the absolute best possible option. The idea can then be formalized by introducing an aspiration level for the agent, written $k \leq \max_{\gamma \in \mathcal{A}} u_{\gamma}$, such that any option α with $u_{\alpha} \geq k$ is satisficing and can therefore be picked [63]. In many cases, the solution is non-unique as a result, which we will see is in fact a very natural setting when considering the solution landscape of many optimization problems. As an illustration, the previously discussed logit rule can be seen as a *satisficing* solution to the bounded rationality modeling problem from our perspective.

1.2.4 Simplicity versus fidelity

Having introduced utility theory as well as choice theory, we should be all set to start designing models, perform numerical simulations, and, soon enough, attempt to describe these models analytically. Before doing so, however, there remains some important points to address in terms of our philosophy and approach.

Throughout the years, the evolution of agent-based modeling has gone hand in hand with the increase in computational power available to researchers, and more recently with the availability of an incredible wealth of data from a variety of sources. An immediate temptation is therefore to complexify models as much as possible, introducing a very large number of parameters and agents, with the idea of calibrating models on real data and to then conduct numerical experiments on an *in silico* version of the world. As an extreme example, some authors have for instance recently proposed a “1:1” model for the Hungarian housing market [65], comprising 4 million Hungarian households built using empirical micro-data, a

housing and rental market, construction sector, banking sector and even demographic trends. At the other end of the spectrum, what is considered to be a very simple model of an economy, the “Mark-0” model of Gualdi *et al.*, still includes no less than 14 adjustable parameters in its barest form [66]. While essential to study large scale economic systems and a necessary effort to eventually improve the forecasts produced by economic theory, this sort of high fidelity agent-based modeling is not the aim of the work presented in this thesis.

In an authentic statistical physics approach, our objective is rather to first construct minimal models presenting a rich *phenomenology*, but with no immediate aspiration at exactly replicating real systems. The objective is to identify how a restricted set of conceptually important ingredients, and the interplay between them, can affect the qualitative behavior of the system. The identification of these ingredients should then help the design of more complete models down the road. In physical modeling, these minimal phenomenological descriptions are often referred to as *toy models*, and have proved absolutely essential to understand the mechanisms behind some incredibly non-trivial observations.

1.3 Toy stories

Having established that we do not aim at studying incredibly detailed agent-based model, let us now illustrate what we mean by a *toy model* in the context of socioeconomic modeling. In the following, we will present two celebrated examples that do not involve the theory of disordered systems but make use of classical statistical mechanics.

1.3.1 The Schelling model, micromotives and macrobehaviours

The Sakoda-Schelling model is an interesting illustration of the role statistical mechanics may play in understanding the emergence of unexpected and complex phenomena in socioeconomics. In the very first issue of the *Journal of Mathematical Sociology* published in 1971, James M. Sakoda proposed a very simple model for the occupation of discrete sites on a two-dimensional square lattice by two distinct populations [67], thereby pioneering agent based models and their application to sociology. In the very following issue of this then novel journal⁷, Thomas Schelling introduced a very similar set of rules, with the idea that some aspects of urban segregation can be explained only by some very mild preference towards being surrounded by neighbors from the same group – skin color in this

⁷To be perfectly precise, the first mention of Sakoda’s model can be traced back to his unpublished PhD thesis completed in 1949, while Schelling’s work can be found in a 1969 working paper [68]. In any case, there is no reason to believe either author took inspiration from the other, the objective of the papers being clearly quite different.




Figure 1.2: Illustration of Grauwin *et al.*’s version of the Sakoda-Schelling model. (a) Utility function of the density peaked at $\rho^* = 0.5$. (b) Initial random configuration at an average density $\rho_0 = 0.5$. (c) Final state with $\rho_q \approx 0.65 > \rho^*$ (shown in (a)) in occupied neighborhoods and presenting some fully empty blocks. Simulation performed with $H = 100$ sites per neighborhood, $Q = 36$ blocks, and $\beta = 10$, courtesy of Max Knicker.

context – as oneself [69]. Unfortunately for Sakoda, history favored Schelling after whom the model was named (and who was eventually awarded the Nobel Memorial Prize in Economic Sciences), while his contribution did not have as big of an impact, having been cited only twice in the ten years following publication.

To be fair, while his model was in fact not very well suited to explain segregation, which is intimately related to past and present public policies rather than self-organization [70, 71], Schelling identified a key point that Sakoda did not. Indeed, the important result is not so much that agents aggregate if they are sectarian, but rather that they aggregate even if their ideal neighborhood is comprised of only slightly more people of their kind. In other words, every agent strictly choosing their own individual preference at each step seems to lead the system to a sub-optimal outcome for everyone – Adam Smith’s “invisible hand” completely fails. Schelling coined this sort of apparent paradox the opposition between *micromotives* and *macrobehaviours*.

Many different versions of the Sakoda-Schelling model can be considered and have their own specificity, as we will see in Chap. 7. In its simplest form, one can start with the occupation problem, i.e. the organization of a single type of agents on the lattice, where agents now have a density-dependent utility. A major simplification proposed by Grauwin *et al.* [14] is then to consider that the lattice is divided in a fixed number Q of neighborhoods, in which we will find a number $n_q \leq H$, $q = 1, \dots, Q$, of the N (fixed) agents in the system. Each site can be occupied by at most one agent, and the total number of site in the system $H \times Q$ is therefore greater or equal to the number of agents. For concreteness, one can take the city to be a square grid of size $L \times L$ with square neighborhoods, as illustrated in Fig. 1.2.

The Sakoda-Schelling model adapted to this geometry can then be described

Chapter 1. Introduction

as follows. At each step of the simulation, an agent is selected at random, and offered a move to an empty site in another neighborhood. The agent will then accept or reject the move based on the difference of utility between the two sites. The important simplification provided by the fixed neighborhoods is that the utility function quantifying the agent's preference, assumed to be the same for all agents, is a function of the block's density $\rho_q = n_q/H$ only. An important feature to recover Schelling's apparent contradiction between micromotives and macrobehaviours is then that this utility should be peaked at a value $0 < \rho^* < 1$, representing an individual's ideal environment, neither too empty nor too full. An example of such a utility function with $\rho^* = 1/2$ is shown in Fig. 1.2(a).

We assume the agents make their decision to accept or reject a proposed move based on the logit rule introduced above. When the agents are poor optimizers and β is small, we observe a uniform distribution of the population among all sites, which will be referred to as the homogeneous phase. In this regime, the average neighborhood density is simply the density of agents in the system $\rho_0 = N/(Q \times H)$. For large β on the other hand, where moves overwhelmingly improve the agent's own satisfaction, we recover the peculiar behaviour described by Schelling and shown in Fig. 1.2: the population concentrates in a few neighborhood where the density will vastly exceed the optimal ρ^* , and that even if we take $\rho_0 = \rho^*$.

Predicting what values the block densities ρ_q will take is not easy at first glance, and is where statistical mechanics can play a part. Indeed, as will be discussed more technically throughout this thesis, taking the appropriate choice for the precise decision rule (sometimes) allows one to derive an effective free energy that is minimized by the prescribed dynamics. Without going into the details just yet, Grauwin *et al.* have shown that the neighborhood free energy reads

$$f(\rho) = \int_0^\rho d\rho' u(\rho') - \frac{1}{\beta} [\rho \log \rho + (1 - \rho) \log(1 - \rho)], \quad (1.7)$$

where $u(\rho)$ is the utility function of an agent and the second term is an entropic contribution that will vanish at $\beta \rightarrow \infty$. Maximizing this quantity under the constraint that the total average density is fixed at ρ_0 (or equivalently performing a double tangent construction on this free energy) is then clearly not the same as maximizing $u(\rho)$, even when the entropy is negligible ($\beta \gg 1$). With these results in mind, the model can then be expanded and enriched. Including a second type of altruistic agents, that do maximize the average utility instead of their own, can for instance be shown to have a strong catalytic effects on the total utility of the system [72].

While perhaps a bit vague at this stage, this simplified Schelling model is an example of the motivation behind this thesis and the idea that statistical physics may be relevant to understand some aspects of socioeconomic models. A puzzle, initially uncovered based on a purely socioeconomic motivation, can be understood

and formalized using the elementary concept of free energy minimization applied to an extremely simplified version of the problem. As we will later see, such a direct mapping between a socioeconomic model and equilibrium statistical mechanics is the exception rather than the rule, but the underlying philosophy of going from the microscopic constituents to the aggregate behaviour will remain the common thread.

1.3.2 The Random-Field Ising Model

In the Sakoda-Schelling model, statistical mechanics is used as a tool to understand and analyze an idealization devised in the socioeconomic context, providing an exact solution to an agent-based model that aims at describing an idealized version of the problem. Often times, the situation is somewhat reversed, however, and results from general physical models can be used to design agent-based models. In other words, while in the Sakoda-Schelling case statistical physics provides a way to understand the passage from micromotives to macrobehaviours, we will now see that the macroscopic behaviour of a physical system displaying features that may be relevant in socioeconomics can provide ideas in the microscopic description of agents and even be interpreted as such.

Originally devised by Imry & Ma in 1975⁸ [74] to understand the influence of random impurities on the properties of magnetic alloys, the Random-Field Ising Model (RFIM) is – as its name conveniently suggests – a generalization of the previously introduced Ising model. As before, it considers N binary spins $S_i = \pm 1$, $i = 1, \dots, N$, which are now all interacting with a random magnetic field that takes a different value at each site. The Ising Hamiltonian therefore includes a new contribution,

$$\mathcal{H} = - \sum_{(i,j)} J_{ij} S_i S_j - \sum_i (H + h_i) S_i, \quad (1.8)$$

with h_i the random component of the field's influence at site i , assumed to be independently drawn from the same zero-mean probability distribution $p(h)$, and H the average magnetic field strength acting on the sample. Now, suppose that a zero temperature system is equilibrated in some configuration and that we vary H . Each spin will then flip only if the local field acting on it changes sign, in other words

$$S_i = \text{sign} \left(\sum_{j \neq i} J_{ij} S_j + h_i + H \right). \quad (1.9)$$

⁸To be perfectly precise, Imry & Ma appear to be the first to have studied the effect of a random field on the ferromagnetic phase, but considered *soft* (continuous) spins and a field-theoretic formulation. In the discrete Ising spin case introduced above, the first description of the model appears to be due to Schneider & Pytte [73] in 1977.




Figure 1.3: Analytical solutions to the RFIM for Gaussian distributed random fields with standard deviation for $J = 1$, (a) $\sigma = 1$, (b) $\sigma = 0.1$, the continuous line showing the solution propagated by increasing H and the dashed line by decreasing H . Insets: graphical resolution of Eq. (1.11) at $H = 0$, dotted and continuous lines representing the left hand and right hand sides respectively, dark and light circles showing stable and unstable solutions respectively.

Clearly, if the random h_i remain constant and are of zero mean, if H is varied continuously then all spin flips will not occur at the same point.

To make a relatively long story short, the model can be solved exactly and rather straightforwardly for $N \rightarrow \infty$ if the network is fully connected and couplings are taken to be constant, $J_{ij} = J/N$ [75]. Characterizing the sample by its total magnetization

$$M = \frac{1}{N} \sum_i S_i, \quad (1.10)$$

one can show that this macroscopic property should satisfy the self-consistent equation

$$M = 2P_>(-H - JM) - 1, \quad (1.11)$$

with the survival function of the random field distribution $P_>(h) = \int_h^\infty ds p(s)$. For any probability distribution, the survival function $P_>(h)$ is a strictly decreasing function satisfying $\lim_{h \rightarrow -\infty} P_>(h) = 1$ and $\lim_{h \rightarrow \infty} P_>(h) = 0$ such that the right hand side of the above equation is strictly increasing in M and goes from -1 to 1. Solving the equation graphically, it immediately appears that it goes from having single solution to three solutions at a critical point (when both sides become tangent at their intersection).

The solution to the equation as well as this possible multiplicity of solutions is illustrated in Fig. 1.3 for $J = 1$ and a Gaussian random field characterized by its standard deviation σ . If the standard deviation is large enough, and the random field is widely distributed, the self-consistent equation has a unique solution and varying H from $-\infty$ to $+\infty$ gives the same solution as varying it in the other direction. When the h_i are tightly distributed (or, fixing σ , if J is increased),

then two new solutions appear and the direction in which H is varied becomes important. As displayed in Fig. 1.3(b), the system displays hysteresis, which can also be confirmed by performing numerical simulations at large N .

Perhaps more interestingly, the switch from $M = -1$ to $M = 1$ can be understood in terms of avalanche dynamics: going from negative to positive H , the first spin flip will be for the site with the largest value of h_i , resulting in a slight increase in the value of M , which in turn will allow the next site with the largest h_i to switch etc. Close to the critical value at which hysteresis comes into play, the precise distribution of avalanche sizes can be computed [75], revealing a power law decay for large avalanches with exponent $-3/2$ (meaning the typical avalanche size diverges in the vicinity of the critical point).

So how is this all related to socioeconomics? Superficially, the abrupt transition observed in the RFIM is already an interesting feature, as many socioeconomic phenomena display this sort of sharp switches, from economic crises to fads and trends. In this context, the model is interesting in the sense that it clearly underlines the trade-off between heterogeneity (governed by the width of the distribution of the h_i) and the tendency to align with the average (here materialized by the term JM). Most importantly, the phenomenology of the RFIM appears to be relevant in the context of widespread changes in opinions, which we can also expect to spread throughout a group in an avalanche-like fashion like spin flips here. As a matter of fact, Michard & Bouchaud [76] have shown that, beyond this *qualitative* appeal, the anomalous exponent relating the slope of the magnetization with H close to the critical point is *quantitatively* compatible with a large number of real world systems. In this work, the authors have indeed collected and measured the scaling relation between the height and the width of the peak in the speed of change of birth rates, sales of cellphones and decay of applause – all somewhat driven by social pressure – and found them to fall within the range of the mean-field prediction of the RFIM.

In light of this empirical evidence, the original fixed point equation of the Random Field Ising Model in Eq. (1.9) can be reinterpreted. Suppose S_i is now one of two options, for example to vote for or against Brexit, to pay or avoid taxes, to buy a new product or refrain etc. The first term is evidently one of imitation, where $\{J_{ij}\}$ is now an interaction matrix that could in principle include both positive and negative entries. The random field h_i can be an individual preconception or preference, while finally the (possibly time-varying) H common to all agents can be interpreted as some public information (e.g. news). The original physical system can then be enriched in this context, for instance through the introduction of non-reciprocity in the interactions or of an underlying network structure. This type of approach, whereby a general model from physics is extended and further completed by ideas tailored to the socioeconomic system it now aspires to represent will be central in this work. In the same way that statistical physics may be

useful to understand socioeconomic models, models from statistical physics can perhaps be of use to provide inspiration in order to recover empirically observable phenomena.

1.3.3 A word on universality

At this stage, the reader uninitiated to statistical mechanics may well wonder why we are even trying to model individuals as imaginary spins or a city as a rectangular grid. The RFIM's scaling exponent may perhaps have been found to approximately match a handful of real world situation, but after all what if this is just a lucky guess? This is actually a very common criticism to the econophysics or sociophysics approaches – how could one even believe they are modeling a socioeconomic systems without making experimental measurements, polls, and data analyses? *A priori*, these toy models, as their name evidently indicate, precisely appear to be mere toys for physicists to play with.

While we won't argue with the necessity of experimental measurements and data analyses, which are at the heart of the physics-inspired approach, there is some justification for the appeal of toy modeling beyond analytical convenience. As mentioned in the introductory remarks, the discovery of *universality* in emergent phenomena has given much more credence to toy modeling for the understanding of even some of the most complex phenomena.

Going back to good old fashioned laboratory experiments, a highly complex phenomenon in fluids which has long intrigued physicists is the so called liquid-gas transition. For sufficiently high temperatures, there is a single phase region, where one can go from a gas to a liquid simply by increasing the density. Below a fluid-dependent critical temperature T_c , however, one observes the emergence of a two-phase region separating exclusively gas and exclusively liquid regions. Now, looking more precisely at the vicinity of this continuous transition, the experimental study of sulphurhexafluoride ($T_c \approx 319$ K) reveals the power law behavior

$$|\rho_+ - \rho_-| \propto |T - T_c|^{0.327 \pm 0.006}, \quad (1.12)$$

where $\rho_{\pm}$ corresponds to the coexisting densities in the two-phase region. Remarkably, repeating the experiment on helium 3 atoms ($T_c \approx 3$ K) yields,

$$|\rho_+ - \rho_-| \propto |T - T_c|^{0.321 \pm 0.006}, \quad (1.13)$$

i.e. an almost undistinguishable value of the exponent despite the difference in T_c [77]. Similar values have also been reported in binary fluids, such as solutions of isobutyric acid and water or carbon disulfide and nitromethane [78].

This apparent universality in the coexistence densities in the vicinity of the critical temperature is very interesting in its own right, but this is not the end of the

story. Suppose we now consider the three-dimensional antiferromagnet DyAlO_3 under a vanishing magnetic field. It is well documented that the material will lose its magnetization if it is heated to a sufficiently high temperature. Approaching the critical temperature $T_c \approx 3.5$ K at which the magnetization M continuously vanishes from the low temperature region, experimental measurements point to a power law behavior of the form [77]

$$M \propto |T - T_c|^{0.311 \pm 0.005}. \quad (1.14)$$

Now truly surprisingly, the exponent is within the range of what was found in the liquid-gas transition!

The introduction of the renormalization group (RG) provides a theoretical framework to understand the existence of these universality classes in which the critical behavior is virtually identical for widely different systems. Without going into the details, the RG demonstrates that the microscopic details of a model typically get “washed out” at the critical point, where everything becomes scale invariant. Beyond this pictorial view, the RG also provides some powerful theoretical predictions. In the case of the three-dimensional ferromagnet described by the Ising model for instance, the RG result for the magnetization exponent is 0.325 ± 0.0015 , which is in very good agreement with experimental results [79].

Another perhaps more formal illustration of universality is the ubiquitous central limit theorem (CLT), which we will invoke at various stages of the thesis. In its simplest form, the theorem states that the arithmetic mean of N independent and identically distributed random variables with finite variance will converge to its mean value (as stated by the law of large numbers) plus *Gaussian* random fluctuations, the variance of which decays as $1/N$. Not only is this true regardless of the underlying distribution of the summed random variables, it remarkably turns out that the CLT holds under much weaker conditions. The random variables can indeed in fact be weakly correlated or not quite identically distributed and the fluctuations will remain Gaussian distributed. As a result, an incredibly wide range of real-world statistics are effectively Gaussian, such as the incidence rate of accidents, the gender-specific distributions of heights and weights, or the velocity of stars [80]. Furthermore, when the variance diverges, a generalized central limit theorem applies, by which the limit distribution of the arithmetic mean is now a Lévy distribution. As a side note, these central limit theorems can also be understood with the renormalization group mentioned above [81]. For a more complete discussion of universality from a mathematical standpoint, see Ref. [80].

To close this digression, universality is a key concept that is at the heart of the effectiveness of statistical mechanics. As the macroscopic features turn out to be almost independent of the microscopic details, universality notably allows to make theoretical predictions using vastly simplified models, which is an extremely powerful idea in the socioeconomic context. Given the undeniable complexity of

human behavior, it seems that our only hope to one day have robust predictions, or at the very least scenarios, describing collective behavior in fact lies in this idea that aggregation erases some of the details. Having said this, all the while avoiding the now proverbial phrase *more is different*⁹, we may embark in our journey at the crossroads between socioeconomics, spin-glasses and out-of-equilibrium dynamics.

1.4 Thesis layout

This thesis is organized as follows. In Chapter 2, we recall some basics of statistical mechanics and key spin-glass results that will be important in the understanding of the different problems we will study, and close the introductory part of the thesis.

We start by exploring two specific socioeconomic problems illustrating the concepts of radical complexity and bounded rationality in Part II. Chapter 3, which is largely based on the published work [1], explores a simple constrained portfolio optimization problem that individual agents could face. We show that the problem has a very large number of quasi-degenerate solutions, challenging the very notion of rationality and common information among agents. Having established the motivation for boundedly rational decision making in the modeling of individual agents, we revisit a classic object of consumer choice theory in Chapter 7, the Slutsky matrix. This disorder-free digression, the results of which have been presented in [2], allows us to illustrate how bounded rationality may be compatible with established empirical results that were thought to support the infinitely rational assumption. Moreover, this practical problem illustrates the possible absence of a global utility (and therefore a global free energy) when agents are selfish and interacting, underlining the necessity for a better understanding of out-of-equilibrium dynamics.

Having well motivated the different ingredients that appear necessary to construct interesting toy models, we enter Part III of the manuscript, dedicated to the “SK-game”, a unifying binary decision model with bounded rationality, non-reciprocity, and learning. The model is introduced in Chapter 5, and its most salient features are exposed using numerical simulations. In Chapter 6, we undertake a more technical analysis of the model. Its fixed point and limit cycle complexity are discussed and we notably employ Dynamical Mean-Field Theory to study the $N \rightarrow \infty$ dynamics of the problem. The results of both the chapters of this part should be included in the forthcoming publication [3].

Part IV is devoted to disorder-free out-of-equilibrium systems – although some aspects remain relevant in the understanding of disordered systems. Chapter 7

⁹Although it is tempting, it would feel very wrong not to cite P.W. Anderson’s illustrious article anywhere in an interdisciplinary statistical physics thesis [82].

corresponds to the results of the collaboration with R. Zakine and A.-C. Becharat published in [5]. In this chapter, we revisit the statistical mechanical treatment of the Schelling model presented above. By removing the necessity for fixed neighborhood and keeping the decision rule as general as possible, we study the robustness of the condensation transition out-of-equilibrium, and exploit links with recent results in the theory of active matter. Chapter 8, which includes and leverages some of the results published with A. Dechant and S.-i. Sasa in [4], finally studies the effect of out-of-equilibrium currents on the relaxation to steady-states that preserve a given Gibbs-Boltzmann distribution, allowing us to isolate the effect of irreversible contributions on the dynamics. While the chapter focuses on a very simple single particle setup, such dynamics may also turn out to be beneficial to accelerate the sampling of complex equilibrium distributions, which could include some disordered agent-based models.

Part V of the manuscript finally discusses future directions and closing remarks.

Key results and messages will be summarized at the end of each chapter, while some of the detailed calculations will be given in dedicated appendices at the end of the manuscript.

Chapter 1. Introduction

Chapter 2

Theoretical foundations

Merely quantitative differences, beyond a certain point, pass into qualitative changes.

Karl Marx

Before jumping into specific socioeconomic problems and their description with statistical mechanics, let us provide some of the theoretical background that will be useful in all subsequent chapters. The goal of this chapter is not to be a self-contained lecture of all aspects of statistical physics, but rather to introduce essential concepts and provide a non-exhaustive summary of some of the most important results in the theory of disordered systems.

2.1 Equilibrium statistical mechanics

2.1.1 Thermodynamics and the Gibbs-Boltzmann measure

When considering atoms or molecules, a natural approach to thermodynamics is to think in terms of the combinatorics of “microscopic” configurations $\{\mathbf{x}\}$. Given the discrete set of possible states of an isolated system, one can indeed attempt at enumerating the number of such states $\Omega(E)$ for which the system will have an energy E . The fundamental postulate of thermodynamics is that all the microscopic states at a certain energy are equally likely. This uniform probability distribution is known as the *microcanonical ensemble*.

$$P(\mathbf{x}|E(\mathbf{x}) = E) = \frac{1}{\Omega(E)}. \quad (2.1)$$

Now, recall that we usually draw an analogy between energy and utility. The elementary postulate of the microcanonical ensemble therefore corresponds to assuming that all choices providing a given level of utility are equiprobable, which

Chapter 2. Theoretical foundations

appears to be reasonable. This being said, the fact that the energy or the utility is strictly *fixed*, as it was assumed that the system is isolated, appears as a fundamental limitation. Beyond the fully rational case where we would assume that the utility level is set to its largest possible amount, any notion of bounded rationality implies that there is a range of possible utility levels rather than a unique value.

Luckily, the necessity to allow for energy fluctuations also quickly arose in physics. Indeed, it is very unlikely for a system to be truly isolated. In most practical cases, the small system that one considers is in fact in contact with a much larger environment. Concretely, if we are interested in studying the properties of a glass of water sitting on a table, it is clear that the conditions of the room in which the experiment takes place will be of paramount importance. In the thermodynamics jargon, the room acts as a *thermostat* for the glass of water: given it is much larger, we can assume that the room is unaffected by what goes on in the glass, and we can therefore take its temperature to be constant throughout the experiment. The room then acts as an infinite thermal energy reservoir to fuel jumps in between microscopic states of different energies, while the *average* internal energy inside of the glass remains constant once it has reached equilibrium. We will shortly see that there is a very natural interpretation for the temperature in an agent-based framework.

The number of combined microscopic states at a given *total* energy level $E' = E_R + E$, where E_R is the energy of the “reservoir” (the room), can be written as

$$\Omega(E') = \sum_{\{\mathbf{x}\}} \Omega_R(E' - E(\mathbf{x})) = \sum_{\{\mathbf{x}\}} \exp(S_R(E' - E(\mathbf{x}))), \quad (2.2)$$

where the number of states of the reservoir $\Omega_R(E_R)$ at a given energy E_R is related to their entropy $S_R(E_R)$ by the microcanonical definition of the entropy

$$S(E) = \log \Omega(E). \quad (2.3)$$

As detailed in any good statistical mechanics textbook, e.g. Ref. [83], one can then use the fact that $E' \gg E(\mathbf{x})$, by virtue of R being a reservoir, to Taylor expand the entropy. Identifying the temperature, which in thermodynamics is formally defined¹⁰ as $T^{-1} = \partial_E S(E)$, the probability to find the system in the state $\mathbf{x}$ is finally given by the celebrated Gibbs-Boltzmann distribution

$$P(\mathbf{x}) = \frac{1}{Z} e^{-\beta E(\mathbf{x})}, \quad (2.4)$$

where $\beta = T^{-1}$ is the inverse temperature and Z is the partition function,

$$Z = \sum_{\{\mathbf{x}\}} e^{-\beta E(\mathbf{x})}. \quad (2.5)$$

¹⁰Note that we have implicitly taken $k_B = 1$ in this discussion, meaning the temperature is taken to be in units of energy.

Remarkably, all terms are independent on the details of the reservoir, whose only remaining property is the temperature.

Computing the logarithm of the partition function then gives access to virtually all thermodynamic quantities of interest, for example the average internal energy

$$U = \langle E \rangle = -\frac{\partial}{\partial \beta} \log Z. \quad (2.6)$$

Unless indicated otherwise, angular brackets will refer to averages with respect to thermal fluctuations. In the following, an equilibrium distribution will refer to the Gibbs-Boltzmann measure for a known energy function of the degrees of freedom.

Now, recall the logit decision rule presented in the introduction, which assigned to the choice $\alpha \in \mathcal{A}$ the probability

$$P_\alpha = \frac{e^{\beta u_\alpha}}{\sum_{\gamma \in \mathcal{A}} e^{\beta u_\gamma}}.$$

Clearly, this is none other than the Gibbs-Boltzmann distribution over the choices of the agent. As a result, in this single-agent noisy decision-making process, the results of standard thermodynamics can be immediately applied through the analogy between energy and utility and by identifying the inverse temperature as the intensity of choice. For example, the relation between the partition function and the internal energy can be directly employed to compute the expected utility of the agent. This being said, and as will be repeatedly stressed at various stages, this will not necessarily be true when there are several agents. Note also that, as discussed in the first chapter, this is only one possible decision rule. Unlike in classical thermodynamics, there is no notion of a thermostat or thermal energy to constrain the passage from the microcanonical to the canonical ensemble.

For completeness, let us mention the last thermodynamic ensemble. In general, any of the ensembles can be seen as the probability distribution maximizing the entropy¹¹ (recall Sec. 1.2.3) under a different constraint [84]. In the microcanonical ensemble, the energy is strictly fixed. In the canonical ensemble, we relaxed this constraint and only fix the energy on average. In the remaining *grand canonical* ensemble, it is the number of particles N that is only fixed on average. Just as the thermostat acts as an unlimited reservoir for energy fluctuations in the canonical ensemble, we imagine that there is an infinite reservoir of particles in the grand canonical ensemble. This particle reservoir is given a *chemical potential* analogous to the temperature. More generally, we will call grand canonical ensemble any

¹¹Here we mean the “statistical” definition of the entropy found by Gibbs and rediscovered by Shannon $S = -\sum_{\{\mathbf{x}\}} p(\mathbf{x}) \log p(\mathbf{x})$, which can be neatly derived from the microcanonical definition of Eq. (2.3) using Stirling’s approximation of factorials and the fact that the entropy is additive [84].

Chapter 2. Theoretical foundations

probability distribution where a constraint (that is not the energy or analogous) is relaxed and only enforced on average instead of being strictly fixed, as will be the case in Chap. 4. In this example, we will see that enforcing a budget constraint at the average level greatly simplifies calculation.

Note that in classical thermodynamics, all ensembles turn out to give identical results in the *thermodynamic limit* – that is when the number of particles N tends to infinity. A quick way to see the equivalence between the microcanonical and canonical ensembles is to generalize Eq. (2.6) to higher order cumulants of the energy, thereby also highlighting the role of the logarithm of Z as a generating function. As it is usually an extensive quantity, $\log Z \sim N$, one can show

$$\frac{\sqrt{\langle E^2 \rangle - \langle E \rangle^2}}{\langle E \rangle} \sim \frac{1}{\sqrt{N}} \xrightarrow{N \rightarrow \infty} 0, \quad (2.7)$$

and similarly for higher order cumulants, meaning that the distribution of E indeed concentrates to an infinitely narrow peak about $\langle E \rangle$. This sort of *concentration of measure* about the average is at the heart of the effectiveness of statistical mechanics to describe high dimensional systems, at odds with the famous “curse of dimensionality”.

Let us close this modest primer of statistical mechanics by mentioning the notion of a *typical* solution. To illustrate this idea, we can write the partition function as a sum over all possible values of the energy instead of all possible states,

$$Z = \sum_{\{E'\}} \Omega(E') e^{-\beta E'} = \sum_{\{E'\}} e^{-\beta(E' - TS(E'))}, \quad (2.8)$$

Where one can identify the Helmholtz free energy

$$F = E - TS(E). \quad (2.9)$$

Now, what we mean by typical is the most probable value (i.e. the *mode* in a more mathematical jargon). The typical value of the energy E^* is the solution to

$$\frac{\partial}{\partial E'} \left(\Omega(E') e^{-\beta E'} \right) = 0, \quad (2.10)$$

which dominates the sum due to the expected exponential nature of both factors. As a result, we can loosely approximate the partition function using a discrete version of Laplace’s method,

$$\log Z \approx -\beta F(E^*). \quad (2.11)$$

In the canonical ensemble, we can use the fact that the standard deviation of E becomes negligible with respect to its mean in the thermodynamic limit to

immediately conclude that the typical and mean values of the energy coincide. In disordered systems, where one has to average over random bonds for instance, Eq. (2.11) shows that the representative quantity to consider will be the average of $\log Z$ and not of Z itself in order to recover the properties of a typical sample.

2.1.2 Detailed balance

So far, we have written the equilibrium distribution describing a statistical mechanical system with a prescribed energy at a given temperature. Remarkably, the only requirement on the underlying system's dynamics is through the fundamental postulate that all microscopic states of equal energy are equiprobable when it is isolated. This is all well and good, but it is legitimate at this point to wonder how a system would practically evolve towards an equilibrium state from a random initial condition. Besides, as we expect a large number of degrees of freedom, the partition function may be difficult to work out analytically, in which case numerical simulations correctly sampling the equilibrium distribution become essential to measure observables of interest.

Perhaps the simplest way to prescribe dynamics given a target probability distribution is to write a Markov chain. Starting from a given initial condition $\mathbf{X}_0$, the idea is to successively jump between configurations $\mathbf{X}_0, \dots, \mathbf{X}_{t-1}, \mathbf{X}_t$ with rates

$$W(\mathbf{x} \rightarrow \mathbf{x}') = \mathbb{P}(\mathbf{X}_{t+1} = \mathbf{x}' | \mathbf{X}_t = \mathbf{x}). \quad (2.12)$$

The assumption that the rate only depends on the current state and not on previous history is what defines the Markovian nature of the dynamics. The associated evolution of the time-dependent probability distribution of the states then follows the master equation

$$\partial_t P(\mathbf{x}, t) = \sum_{\mathbf{x}' \neq \mathbf{x}} [P(\mathbf{x}', t)W(\mathbf{x}' \rightarrow \mathbf{x}) - P(\mathbf{x}, t)W(\mathbf{x} \rightarrow \mathbf{x}')], \quad (2.13)$$

which essentially enforces the conservation of the total probability [85]. In the steady-state, we obviously require $\partial_t P(\mathbf{x}, t) = 0$. A *sufficient* condition for this steady-state to sample the target distribution $P(\mathbf{x})$ is therefore the so-called *detailed balance* criterion

$$P(\mathbf{x})W(\mathbf{x} \rightarrow \mathbf{x}') = P(\mathbf{x}')W(\mathbf{x}' \rightarrow \mathbf{x}), \quad (2.14)$$

enforcing the balance between outgoing and incoming probability flux at the “microscopic” level of each state.

Beyond its apparent simplicity, the detailed balance condition carries rather profound implications. Indeed, it is also the only choice which enforces *time reversal symmetry* (TRS) of the dynamics: the forward and backward trajectory

Chapter 2. Theoretical foundations

in time between two states are equiprobable. In fact, the precise definition of an equilibrium steady-state (in opposition with a nonequilibrium steady-state) is that it precisely satisfies TRS¹². In the rest of the thesis, we will therefore use detailed balance violating and irreversible interchangeably.

As will be discussed in detail in Chap. 8, one can sample the correct Gibbs-Boltzmann steady-state distribution while violating TRS by finding transition rates W that satisfy *global balance* (i.e. setting the rhs of Eq. (2.13) to zero) instead of detailed balance. In this case, the system is out-of-equilibrium, even though its steady-state distribution corresponds to that of a known equilibrium system.

Suppose we place ourselves in the equilibrium context and therefore want to sample the Gibbs-Boltzmann distribution while satisfying detailed balance. Plugging in the expression of $P(\mathbf{x})$ in Eq. (2.14), the rates must satisfy

$$\frac{W(\mathbf{x} \rightarrow \mathbf{x}')}{W(\mathbf{x}' \rightarrow \mathbf{x})} = e^{-\beta \Delta E}, \quad \Delta E = E(\mathbf{x}') - E(\mathbf{x}). \quad (2.15)$$

An essential aspect of these rates is that they do not require any knowledge of the partition function, but only the energy difference between two specific states (which is usually easy to compute).

There are then many choices of transition rates that will satisfy this detailed balance criterion [87]. The most common, which we will for instance use in the simulations of Chap. 4, is the Metropolis algorithm [88]

$$W(\mathbf{x} \rightarrow \mathbf{x}') = \min \left(1, e^{-\beta \Delta E} \right). \quad (2.16)$$

Another common choice for spin systems that is worth mentioning is the Glauber dynamics [89]

$$W(\mathbf{x} \rightarrow \mathbf{x}') = \frac{e^{-\beta \Delta E}}{1 + e^{-\beta \Delta E}}. \quad (2.17)$$

Interestingly, Glauber dynamics are also a rather common way to state the logit decision rule. Indeed, in the socioeconomic context it is in fact more natural to define a model through the transition rate, representing an agent's behavior, rather than some stationary distribution of the choices.

2.1.3 Langevin dynamics

While these transition rates can be readily interpreted in the agent-based context, this is often not the case in physics. As previously mentioned, there are many

¹²A modern definition is that an equilibrium steady-state is one where there is a zero entropy production rate at the trajectory level, a quantity which measures the breaking of TRS [86].

possible choices satisfying detailed balance and it might be awkward to postulate how the system evolves in reality if one is interested in the precise dynamics. Besides, if the state $\mathbf{x}$ is a continuous variable, it might be difficult to correctly sample the entire support of its probability distribution with discrete rates.

A somewhat more physical approach is to follow the continuous overdamped Langevin dynamics

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t)) + \sqrt{2T}\boldsymbol{\xi}(t), \quad (2.18)$$

where $\boldsymbol{\xi}$ is a vector of independent, identically distributed white noises, i.e. satisfying

$$\langle \xi_i(t) \rangle = 0, \quad \langle \xi_i(t) \xi_j(s) \rangle = \delta_{i,j} \delta(t-s). \quad (2.19)$$

Physically, if the degrees of freedom can be understood as a position, this evolution can be interpreted as Newton's second law in the case where the ratio of the mass to the friction goes to zero (e.g. a colloid in a viscous fluid). In the absence of the force field $\mathbf{F}$, the above is for instance simply the Wiener process describing the large time behavior of the pollen particle in water observed by Robert Brown [90,91].

With these overdamped Langevin dynamics, the evolution of the probability density of $\mathbf{x}$ is given by the Fokker-Planck equation, i.e. the continuous analogue of the master equation

$$\partial_t P(\mathbf{x}, t) = -\nabla \cdot [\mathbf{F}(\mathbf{x})P(\mathbf{x}, t) - T\nabla P(\mathbf{x}, t)] = -\nabla \cdot \mathbf{J}(\mathbf{x}, t), \quad (2.20)$$

where $\mathbf{J}$ is the probability current [92]. The equivalent of *global balance* necessarily satisfied in the steady-state is therefore for the probability current to be *divergence-free*. The sufficient condition that is analogous to *detailed balance* is then to have zero probability *current* in the steady-state,

$$\mathbf{F}(\mathbf{x})P(\mathbf{x}) - T\nabla P(\mathbf{x}) = 0 \quad \Leftrightarrow \quad \nabla \log P(\mathbf{x}) = \beta \mathbf{F}(\mathbf{x}), \quad (2.21)$$

i.e. $\mathbf{F}$ must be a conservative vector field and is thus purely gradient (i.e. in three dimensions $\nabla \times \mathbf{F}(\mathbf{x}) = 0 \forall \mathbf{x}$). As in the discrete case, this much stronger condition is necessary for the dynamics to be TRS and therefore for the steady-state to be an equilibrium one. As a consequence, any non-gradient term will drive the system out-of-equilibrium. Once again, out-of-equilibrium dynamics for which the probability current is not zero everywhere may nonetheless lead to the same steady-state as some equilibrium system (Chap. 8).

Replacing $P(\mathbf{x})$ with the equilibrium Gibbs-Boltzmann measure, we immediately find that the required force field is given by $\mathbf{F}(\mathbf{x}) = -\nabla E(\mathbf{x})$, and the overdamped Langevin equation becomes

$$\dot{\mathbf{x}}(t) = -\nabla E(\mathbf{x}(t)) + \sqrt{2T}\boldsymbol{\xi}(t). \quad (2.22)$$

The above simply correspond to a gradient descent with noise minimizing the objective function $E(\mathbf{x})$, providing an appreciable intuitive picture.

2.1.4 A word on the relaxation time

Regardless of the prescribed dynamics, an important aspect to mention is the time required for the probability distribution to reach equilibrium. Formally, this relaxation time can be studied by diagonalizing the transition rate matrix in the discrete case or the Fokker-Planck operator in the continuous case. Asymptotically, the distance between the probability distribution at time t and its stationary limit will decay exponentially as $e^{-t/\tau}$, where the characteristic relaxation time τ is the inverse of the first nonzero eigenvalue¹³ (in absolute value) [93].

At low temperatures, the relaxation time will be dominated by the potentially large energy barriers between local minima of the energy function. This can be understood pictorially with Eq. (2.22): if the noise term is small, it will take a very long time for a large enough fluctuation to lead to the escape a local minimum even though another minima might have a smaller energy. More precisely, the time to escape a valley with a free energy barrier¹⁴ of size ΔF

$$\tau \propto \exp(\beta \Delta F), \quad (2.23)$$

first observed empirically by Arrhenius [96]. As we will see shortly, this scaling is crucial to understand potential failures of classical statistical mechanics.

2.1.5 Ergodicity

The assumption that all states of identical energies are equally likely to be explored by the system is at the heart of the thermodynamics presented in Sec. 2.1.1. This conjecture, also known as the *ergodic hypothesis*, is actually unproven for most systems of interest. It is therefore an essential aspect to keep in mind, particularly when the classical theory fails to replicate some experimental or numerical results.

As a matter of fact, even some relatively simple systems already display quite evident violations of the ergodic hypothesis in the form of *spontaneous symmetry breaking*. The standard Ising ferromagnet is perhaps the most well known example. In its simplest form, its energy is given by the Hamiltonian

$$\mathcal{H} = -\frac{J}{N} \sum_{(i,j)} S_i S_j, \quad (2.24)$$

with $S_i = \pm 1$, $i = 1, \dots, N$, and where (i, j) indicates a sum over bonds, say nearest neighbors on a two-dimensional lattice. (The attentive reader might notice that this is quite conveniently the mean-field Random Field Ising Model of

¹³In the reversible (equilibrium) case, all eigenvalues are real.

¹⁴It is essential to note that the important quantity is the *free energy* difference and not simply the energy difference. As highlighted in [94], one can in fact have $\Delta E < 0$ and still be confined by the entropic term, see also [95] for an illustration of extremely slow relaxation without any energy barriers.

Sec. 1.3.2 without the random field.) Importantly, this Hamiltonian has “ $\mathbb{Z}_2$ symmetry”, meaning that flipping the sign of *all* spins simultaneously does not change the value of the energy. Now, in the low temperature region, the magnetization $M = \frac{1}{N} \sum_i S_i$ will be non-zero. In the absence of an external field as presented so far, the magnetization has two equally likely values, $M = \pm M_S$, but every instance of the system seemingly picks one of the two. The initial symmetry of the Hamiltonian is thus said to be spontaneously broken because the resulting equilibrium state is not invariant under the same symmetry. Note that while spontaneous symmetry breaking implies ergodicity breaking (see below), the converse is not true. As we will see with spin-glasses, there can be true ergodicity breaking without apparent symmetry breaking.

At its root, ergodicity is a dynamical property. Indeed, it implies that the entirety of the configuration space can eventually be explored by the dynamics. Given a sufficiently long horizon, time averaged observables should therefore converge to their ensemble average. Clearly, this is not true when there is spontaneous symmetry breaking: if the system is truly stuck in an equiprobable $\pm M_S$ state, then averaging even over an infinite time will not give a zero mean magnetization predicted by the Gibbs equilibrium state.

To understand this issue, let us dwell on the dynamics of the simple $2d$ -Ising model. Suppose the system is in a state corresponding to the $+M_S$ magnetization. The least costly way to flip all spins and go to a $-M_S$ state is a sweeping line defect across the lattice [94]. It can be argued that the free energy barrier associated to such a defect scales with its length, which is $O(\sqrt{N})$ in two dimensions. As a result, ΔF diverges with N and, from the Eyring-Kramers scaling of Eq. (2.23), the free energy landscape is completely separated in two distinct regions in the thermodynamic limit $N \rightarrow \infty$. As a consequence, one has to be careful in the analytical treatment of the problem.

Such a situation is referred to as *true* ergodicity breaking. In other cases, the configuration space is not strictly separated by infinitely large barriers, yet the time required to explore an infinite system still diverges, and we therefore observe non-trivial out-of-equilibrium dynamics. This is known as *weak* ergodicity breaking, and can for instance occur in so-called “trap” models [97] or with logarithmic potentials [98], as well as some simplified spin-glasses [99] (see Sec. 2.2.5 below).

2.2 Disordered systems

In the introduction chapter, we mentioned the Edwards-Anderson spin-glass, in which neighboring spins tend to align or anti-align based on randomly drawn interactions. Due to its very limited connectivity – $2d$ nearest neighbors in a d -dimensional lattice – the EA model is extremely difficult to describe analytically,

Chapter 2. Theoretical foundations

and as a matter of fact remains unsolved to this day.

A natural and more analytically favorable extension that we also mentioned above is the “mean-field” version proposed by Sherrington and Kirkpatrick [28], for which the Hamiltonian now reads

$$\mathcal{H} = -\frac{1}{2} \sum_{j \neq i} J_{ij} S_i S_j, \quad (2.25)$$

where the random symmetric couplings $J_{ij} = J_{ji}$ are still Gaussian but with a variance that now scales as $1/N$,

$$P(J_{ij}) = \sqrt{\frac{N}{2\pi J^2}} e^{-\frac{N}{2} \left(\frac{J_{ij}}{J}\right)^2}. \quad (2.26)$$

In the following, we will take the SK model as the prototypical example of a disordered system, as the analytical challenges it poses are ideally suited to illustrate essential concepts, which will be mentioned throughout the thesis.

2.2.1 The replica method

In order to solve the model, the standard statistical mechanics path is to compute the average of the logarithm of the partition function to obtain the typical free energy density,

$$f = -\frac{1}{\beta N} \overline{\log Z}, \quad Z = \sum_{\{S\}} e^{-\beta \mathcal{H}(\{S_i\})}. \quad (2.27)$$

Here, overlines indicated averages with respect to the distribution of disorder given in Eq. (2.26). To tackle the average of the logarithm, one can invoke the so-called replica trick

$$\overline{\log Z} = \lim_{n \rightarrow 0} \frac{\overline{Z^n} - 1}{n}, \quad (2.28)$$

where n will be treated as an integer throughout the calculation before being famously (or infamously) analytically continued to 0 [44].

Rewriting the power as the product between n replicas of the system,

$$Z^n = \prod_{a=1}^n \sum_{\{S^a\}} e^{\frac{\beta}{2} \sum_{j \neq i} J_{ij} S_i^a S_j^a}, \quad (2.29)$$

and performing the average on the symmetric Gaussian bonds, we have

$$\overline{Z^n} = \prod_{a=1}^n \sum_{\{S^a\}} \exp \left(\frac{1}{4} \beta^2 N n + \frac{1}{2} \beta^2 N \sum_{a < b} \left(\frac{1}{N} \sum_i S_i^a S_i^b \right)^2 \right). \quad (2.30)$$

Introducing the $n \times n$ overlap matrix

$$Q_{ab} = \frac{1}{N} \sum_i S_i^a S_i^b, \quad a \neq b \quad (2.31)$$

which appears in Eq. (2.30) following the Gaussian average, the averaged replicated partition function can be recasted (after a bit of work) as

$$\overline{Z^n} = \int \prod_{a < b} \left(\sqrt{\frac{N\beta^2}{2\pi}} dQ_{ab} \right) \exp(-NA[\mathbf{Q}]). \quad (2.32)$$

Such an integral can be computed in the $N \rightarrow \infty$ limit with Laplace's method,

$$\overline{Z^n} \underset{N \rightarrow \infty}{\sim} \exp(-NA[\mathbf{Q}^*]), \quad (2.33)$$

where $\mathbf{Q}^*$ is the minimum (or saddle point in the complex case) of the effective action A [100].

A priori, it seems perfectly adequate to treat the replicas as interchangeable, and to assume that the overlap matrix has the much simplified structure

$$Q_{ab} = q \quad \forall a, b, \quad a \neq b. \quad (2.34)$$

Solving the complete set of equations extremizing A within this *ansatz*, this “replica-symmetric” (RS) solution finally yields the relatively simple self-consistent relation

$$q = \int_{-\infty}^{\infty} \frac{dz}{\sqrt{2\pi}} \tanh^2(\beta J \sqrt{q} z) e^{-\frac{1}{2}z^2}, \quad (2.35)$$

where this overlap can be interpreted as the Edwards-Anderson order parameter, $q = \overline{\langle S_i \rangle^2}$ [28].

With this in mind, the solution appears perfectly suitable: we have $q = 0$ for temperatures above $T_c = J$, and $q > 0$ below, reaching $q = 1$ for $T = 0$, in line with the phenomenology of the EA model and with numerical simulations. Looking closer, however, this ideal picture starts to crumble. First, it is easy to observe that the $T \rightarrow 0$ behavior of q does not match numerical experiments, approaching unity exponentially fast rather than quadratically. Calculating the average energy per spin, one also finds $U(T = 0) \approx -0.79$, at odds with numerics that suggest $U(T = 0) \approx -0.76$ for $N \rightarrow \infty$ [101]. While these discrepancies are already sufficient to convince most theoretical physicist that the theory might not be entirely correct, things get worse. As the ultimate *coup de grâce*, the RS solution displays a negative zero-temperature entropy¹⁵, $S(T = 0) \approx -0.17$.

¹⁵In a discrete system, the entropy is non-negative by definition, as it is the logarithm of an integer number of configurations (Eq. (2.3)).

Further investigation confirms the problem with the RS solution below the critical temperature. Looking at the second derivative of the effective action A on which the saddle point approximation is taken, it indeed turns out that this solution is in fact unstable below the temperatures defined by the so-called de Almeida-Thouless (AT) line [102]. In the SK model, the AT line coincides with the critical temperature, and the replica-symmetric solution is invalid in the entire spin-glass phase.

2.2.2 TAP approach

Facing the failure of the replica-symmetric solution and convinced that its root was the order in which the $N \rightarrow \infty$ and $n \rightarrow 0$ limits are taken, Thouless, Anderson and Palmer attempted to construct an alternative approach [103]. Indeed, given the fully connected nature of the model, the mean-field limit should be exact, and the problem well described by the magnetization at each site, $m_i = \langle S_i \rangle$. It would therefore seem effective and natural to take a mean-field description *before* averaging over the disorder distribution.

Brutally replacing the effect of the other spins by their average, the magnetization at site i appears to satisfy

$$m_i = 2\mathbb{P}(S_i = 1) - 1 = \tanh \left(\beta \sum_j J_{ij} m_j \right), \quad (2.36)$$

known as the Naive Mean-Field Equation (NMFE) in the spin-glass literature [104]. Unfortunately, the NMFE is very obviously wrong to describe the SK model, as studying the linear stability of the paramagnetic $m_i = 0$ solution gives $T_c = 2J$ [44], twice the correct value. In other contexts however, we will see that it is relevant and that it shares many of the interesting properties of the correct TAP equations of the SK model (see below).

To obtain these correct TAP equations, it is actually essential to subtract the “reaction” term, corresponding to the influence on the i^{th} spin on all others (which is $O(1/N)$ but acts on $O(N)$ of its colleagues). Proceeding with a “cavity” approach, i.e. considering the effect of an $N + 1^{\text{th}}$ spin on a system of size N , it is possible to show that the correct TAP equations read

$$m_i = \tanh \left(\beta \sum_j J_{ij} m_j - \beta m_i \sum_j J_{ij}^2 (1 - m_j^2) \right). \quad (2.37)$$

From there, one can then define a TAP free energy F_{TAP} such that $\partial_{m_i} F_{\text{TAP}} = 0$ gives the TAP equations, which will be important to understand the relative weight to give to different potential solutions.

Now, these equations can reassuringly be checked to be consistent, in the sense that their careful study for $T \rightarrow 0$ is compatible with q approaching unity quadratically and a positive value of the entropy. However, in the very words of Thouless, Anderson and Palmer, “*finding solutions [...] is not much easier than the original problem*”, as the equations are formulated for a given draw of the disorder.

2.2.3 Complexity

An essential aspect explaining the difficulty of treating these TAP equations is the (very) large number of solutions they admit. Very pictorially, one can imagine the TAP free energy landscape as a high dimensional analogue to Fig. 2.1, where each local minima is a solution to the TAP equations. Indeed, for a given realization of the disorder, the number of solutions $\mathcal{N}_J$ grows exponentially with the size of the system. For a given draw of the disorder, this number can be computed “directly”, by counting the fixed points among all possible configurations. Informally,

$$\mathcal{N}_J = \int \prod_{i=1}^N dm_i \delta(\text{TAP eq.}). \quad (2.38)$$

Note that one must be careful to perform the proper change of variable such that the Dirac δ has a unit norm when the configuration is a solution to the TAP equations [105].

Computing the average of this number of solutions over the disorder, we obtain the “annealed” complexity

$$\Sigma_{\text{ann.}} = \lim_{N \rightarrow \infty} \frac{1}{N} \log \overline{\mathcal{N}_J}. \quad (2.39)$$

However, and as will become clear in the practical case we will consider, the exponential nature of the number of solutions means that this quantity may display very large fluctuations from one realization to another. A more representative measure of the *typical* number of solutions therefore comes in the form of the “quenched” complexity

$$\Sigma_{\text{quen.}} = \lim_{N \rightarrow \infty} \frac{1}{N} \overline{\log \mathcal{N}_J}. \quad (2.40)$$

The disorder average of the logarithm can then be computed using the previously introduced replica trick.

There is actually a fair bit of controversy and still many unknowns surrounding the complexity of the TAP equations at finite temperature. As we will discuss in more detail in Chap. 6, the original annealed solution worked out by Bray & Moore in the 1980s [106] was strongly challenged by the more mathematically




Figure 2.1: Two-dimensional idealization of a complex free energy landscape. A large number of local minima of similar depth are separated by large barriers.

sound “supersymmetric” solution developed in the early 2000s [107] (the potential problem having first been identified by Kurchan in 1991 [108]). The Bray & Moore solution was eventually saved by the fact that TAP states are not stable minima but are in fact saddles with a single unstable direction in the $N \rightarrow \infty$ limit, justifying the broken supersymmetry [109, 110]. Note that in the Bray & Moore solution, quenched and annealed complexities coincide for solutions above a certain free energy threshold [106], although we will see that this is not always true.¹⁶ In any case, the important result to remember is that there exists an extremely large number of solutions to the TAP equations for a given interaction matrix. At $T = 0$, the result of Tanaka and Edwards $\overline{N}_J \approx e^{0.1992N}$ [112] holds and illustrates this reality: for $N = 100$ there are close to half a billion solutions!

Interestingly, the direct approach of enumeration is very general as it does not require one to know what the equilibrium measure is, and can be freely employed for any random equation. It should be noted, however, that an alternative approach relating the complexity to the (replicated) free energy exists when the system is described by a Hamiltonian [113]. In the portfolio problem presented in Chap. 3, we will also see that the annealed complexity can be approximated with a cavity-like approach, albeit in a somewhat less controlled fashion.

¹⁶In general, the annealed average is an upper bound of its quenched counterpart by Jensen’s inequality [111] and the concavity of the log.

2.2.4 Replica symmetry breaking

The study of the TAP equations has shed light on an essential aspect of the spin-glass solution landscape: the explosion of the number of local free energy minima. This multiplicity of solutions in turn has an immediate effect on the problem's thermodynamics, which we can now relate to the failure of the replica symmetric ansatz previously introduced.

When considering TAP states, attempting to measure thermodynamic properties (such as the average energy for instance) by taking an arithmetic average over solutions turns out to be completely wrong. Instead, if one is interested in the observable O , the correct approach is in fact to take

$$\langle O \rangle = \sum_{\alpha} \langle O \rangle_{\alpha} w_{\alpha}, \quad w_{\alpha} \propto e^{-\beta F_{\alpha}}, \quad (2.41)$$

where $\langle O \rangle_{\alpha}$ and F_{α} are the averaged observable and the TAP free energy evaluated in the TAP state α respectively. The states α are then referred to as pure states, the linear combination of which leads to the Gibbs equilibrium state of the system [44]. Of course, given the exponential weight, only a relatively small fraction of the solutions with the smallest free energies will actually contribute to the thermodynamics. What is then of paramount importance is that in fully connected models such as the SK, the energy barriers between the low energy TAP states actually grows with N and therefore diverges in the thermodynamic limit $N \rightarrow \infty$ [114]. Schematically, the free energy landscape is formed of distinct valleys, each corresponding to a TAP state, that are completely separated from one another (just as discussed in Sec. 2.1.5).

To study the problem averaged over the disorder, having a single value of the overlap as in the replica symmetric case is therefore not enough [115]. Instead, the correct approach is to study an overlap *distribution*,

$$P(q) = \overline{\sum_{\alpha, \beta} w_{\alpha} w_{\beta} \delta(q_{\alpha\beta} - q)}, \quad q_{\alpha\beta} = \frac{1}{N} \sum_i m_i^{\alpha} m_i^{\beta}. \quad (2.42)$$

As previously discussed, this is not technically specific to spin-glasses, as an Ising ferromagnet without any external field also has an overlap distribution $P(q) = \frac{1}{2} \delta(q - M_S^2) + \frac{1}{2} \delta(q + M_S^2)$, see Sec. 2.1.5 above. What is specific to the spin-glass case is that the exponential number of quasi-degenerate solution means that $P(q)$ is a continuous function that is highly non-trivial.

The computation of the overlap distribution and the associated Parisi function $q(x)$, defined through the relation $\frac{dx}{dq} = P(q)$, is where Parisi's replica breaking "scheme" comes into play [116, 117]. As this thesis does not contain *bona fide* replica calculations, the procedure will not be detailed here. This being said, the conceptual importance and dynamic consequences of replica symmetry breaking

will be extremely important to understand some of the results presented in the manuscript.

2.2.5 Dynamics

As might be expected at this stage, the extremely peculiar nature of the spin-glass phase and its rugged free energy landscape (picture something like Fig. 2.1 but much more extreme) inevitably has consequences on the problem's dynamics. Simulating the SK model with a discrete Markov chain, using Glauber dynamics for instance, some essential properties quickly appear below the critical temperature [118–120].

Defining the two-point correlation function

$$C(t, t') = \frac{1}{N} \sum_i \overline{\langle S_i(t) S_i(t') \rangle}, \quad (2.43)$$

the most striking feature is probably that, for sufficiently large values of N , it never reaches a time-translation invariant form $C(t, t') = C(t - t')$ as one expects in a finite-temperature equilibrium state. Instead, one observes *aging*, that is that the behavior of $C(t_w, t_w + t)$ remains a function of the waiting time (i.e. the age of the system) t_w , even when it is taken to be large. The correlation function, which will be discussed again in Chap. 6, most often takes the additive form

$$C(t_w, t_w + t) = C_{\text{relax}}(t) + C_{\text{aging}}(t, t_w). \quad (2.44)$$

Such a situation is illustrated in Fig. 2.2. The leftmost part of the curves is approximately invariant with the age of the system, corresponding to the $C_{\text{relax}}(t)$ regime, while the long time behavior is expected to be of the form

$$C_{\text{aging}}(t, t_w) = \sum_i \mathcal{C}_i \left(\frac{h_i(t)}{h_i(t_w)} \right) \quad (2.45)$$

in mean field models, where the (potentially infinite) sum runs over different “time sectors” and the h_i may be nonlinear functions [121].

Eq. (2.44) can be superficially understood by picturing a rugged energy landscape as in Fig. 2.1. At low temperatures, one can imagine that the system remains stuck in local minima for extended periods of time. The relaxation component then represents the first decorrelation which occurs due to thermal noise within a given basin of attraction, and that is therefore independent of the age of the system. The second component, on the other hand, represents the decorrelation following a large thermal fluctuations leading to another valley in the configuration space, with possibly a smaller free energy. As the system gets older, the configuration is likely located deeper and deeper in the free energy landscape. The time required




Figure 2.2: Two-point correlation function in the spherical SK model at $T = 0.3$ for different system ages, taken from [123]. Inset: same quantity represented as a function of the time rescaled by the age, t/t_w , showing a clear collapse in the aging regime.

for a thermal fluctuation to overcome the larger energy barrier therefore increases, inducing a dependency on t_w . This complex relaxation behavior is characteristic of the *weak* ergodicity breaking phenomenon mentioned above [97]. Note that this weak ergodicity breaking may occur in addition to (or in combination with) the *true* ergodicity breaking¹⁷, as in the SK model discussed here [122].

An interesting example of a simplified spin-glass model displaying aging is the *spherical* SK model (SSK). The SSK is essentially a continuous version of the SK model [124]. Instead of Ising spins $S_i = \pm 1$, we now take “soft” spins $\sigma_i \in \mathbb{R}$ satisfying the spherical constraint $\sum_i \sigma_i^2(t) = N \forall t$. Following Eq. (2.22), the overdamped Langevin equation associated to the relaxation of the same SK Hamiltonian in this continuous setting is then given by

$$\dot{\sigma}_i(t) = -u(t)\sigma_i(t) + \sum_j J_{ij}\sigma_j(t) + \sqrt{2T}\xi_i(t), \quad i = 1, \dots, N, \quad (2.46)$$

where $\xi_i(t)$ is a Gaussian white noise, and the confining strength $u(t)$ is adjusted to keep the spherical constraint satisfied at all times. Importantly, this spherical SK model has a significantly simpler free energy landscape relative to its Ising counterpart that we have discussed so far. Indeed, the spherical version admits an exact replica symmetric solution with only two equilibrium states in the low temperature phase [124]. Nonetheless, the full resolution of its dynamics [123] has demonstrated that the model retains interesting and non-trivial aging effects, with $C_{\text{aging}}(t, t_w) = \mathcal{C}\left(\frac{t}{t_w}\right)$ as shown in Fig. 2.2, and therefore displays weak ergodicity breaking. These complex dynamics are inherently related to the fact that the

¹⁷The main marker of true ergodicity breaking is to have $\lim_{t \rightarrow \infty} C(t_w, t_w + t) > 0$ whereas we expect $\lim_{t \rightarrow \infty} C(t_w, t_w + t) = 0$ in the weak case.

system does not reach equilibrium, even in this simpler setting, highlighting the important difference between the statics and the dynamics of these problems.

2.3 Summary of the key notions

Let us summarize and clearly name some of the concepts that have been mentioned in this section and that will be important throughout the thesis:

- When a system is defined by a global energy function, or *Hamiltonian*, it has a Gibbs-Boltzmann measure and its static properties are studied with the logarithm of its partition function (i.e. its free energy). Dynamics are usually prescribed in the form of a detailed balance satisfying Markov Chain or overdamped Langevin equation with only gradient contributions.
 - If it is not disordered, the dynamics recover the statics and all is usually well¹⁸, barring possible spontaneous symmetry breaking. This is for instance the case of the system studied in Chap. 4,
 - If it is disordered, the statics will likely display replica symmetry breaking due to a large number of solutions, as in the problem discussed in Chap. 3. As a consequence, any dynamics will probably get stuck in some region of the solution space. The system displays weak and sometimes strong ergodicity breaking, and an equilibrium state might *never* be reached. In this case, we have a Hamiltonian system with relaxational dynamics that is nonetheless *perpetually out-of-equilibrium*.
- When a system is defined through its dynamics, it may or may not be in equilibrium.
 - If the dynamics are relaxational, i.e. can be rewritten as a detailed balance satisfying Markov Chain or overdamped Langevin equation with only gradient contributions, we are back to the previous situations in which the Gibbs-Boltzmann distribution is directly known,
 - If this is not the case, there will always be some probability currents in the steady-state. This is then a *bona fide* (nonrelaxational) out-of-equilibrium system, for which the steady-state distribution is not known in general, as the one studied in Chap. 7. To find the steady-state distribution, one has to solve the Fokker-Planck or master equation, which might not be tractable analytically. It is possible that this steady-state distribution matches a known equilibrium distribution, which will be the focus of Chap. 8, but this does not mean it is in equilibrium,


¹⁸It is actually possible to have a glassy regime and aging effects without disorder, as demonstrated in [125] for instance – but the main idea stands.

- It may very well be the case that such an out-of-equilibrium system is also disordered, and displays weak or strong ergodicity breaking. This is for example the case for a spin-glass like system with a non-symmetric random interaction matrix, which will be the focus of Chaps. 5 and 6. This is, in a sense, the worse (or best, everything is a matter of perspective) case scenario, as the steady-state distribution is unknown and might never be truly reached due to the glassy dynamics.

In this summary, we have discretely introduced the notion of *relaxational* (and conversely *nonrelaxational*) dynamics. Given its presence in the very title of this manuscript, it may be worth clarifying what exactly this implies. In the context of this thesis, relaxational dynamics will refer to any evolution which minimizes a global free energy-like quantity. As a result, any detailed balance satisfying (i.e. reversible) stochastic process falls within this definition, as hinted above. However, this term also extends to *deterministic* dynamics with a Lyapunov function that is monotonously decreased. Considering the evolution in time of a smooth *field* for instance, as will be done in Chap. 7, dynamics that can be written as the functional derivative of a free energy-like quantity are then also relaxational.

Key takeaways

- The logit rule is just another name for the Gibbs-Boltzmann distribution of equilibrium statistical mechanics over the choices of the agent.
- If the decision rule is applied to a system-wide utility function common to all agents, the dynamics satisfy *detailed balance* and there is an *equilibrium* steady-state that is entirely described by the free energy.
- When the system features random heterogeneous interactions, it is said to be *disordered*.
- Disordered systems, such as the Sherrington-Kirkpatrick spin-glass, are characterized by an extremely complex free energy landscape with an exponential number of metastable states.



- This complexity generically leads to *ergodicity breaking*, meaning the dynamics are trapped in some regions of the solution space and the system truly never reaches its equilibrium state.

Part II

Radical complexity and detailed balance violation in socioeconomics

Chapter 3

Rationality versus complexity: the example of portfolio optimization

Faites vos jeux... Rien ne va plus.

Chevalier de Balibari

As the very motivation for this thesis, we have argued that the “homo economicus” viewpoint traditionally taken in socioeconomics is inherently limiting to understand collective behaviors, and that statistical mechanics may be useful to go beyond the rational representative agent paradigm. Before jumping into the modeling of collective systems with interacting agents, let us first use a simple example to further motivate the necessity for bounded rationality.

In an agent-based model, each individual will make decisions based on some form of utility maximization. At this stage, we have also stressed that spin-glasses exemplify how simple optimization problems may suffer from an extremely large number of solutions, leading the optimization procedure to be “NP-hard”, i.e. quickly computationally untractable. Bringing these two ideas together, it is natural to ask if there are some realistic situations in which the maximization procedure that is expected to be carried out by the agents is simply too difficult and sensitive to be systematically and identically performed by all. As we will see, another important consequence of this complexity is that it does not only challenge the idea of rational decision-making, but also of common information at the heart of game theory.

In this chapter, we reproduce the contents of [1], written under the supervision of J.-P. Bouchaud and M. Benzaquen, barring a slight reorganization and some

minor changes in phrasing and references.

3.1 The portfolio optimization problem

The idea of maximizing gains while minimizing risk for a portfolio of fluctuating assets is one that is both at the origin and at the heart of quantitative finance. As early as 1952, Harry Markowitz derived the general formula for the portfolio with the smallest variance for a desired return [126], assuming the asset fluctuations are fully described by a covariance matrix. Since then, Markowitz portfolio theory has stood as a reference in portfolio management, and questions surrounding optimal portfolios have been a very successful playground for physicists. By adapting calculations from the physics of disordered systems, several theoretical results have been obtained, mainly around the phase transition observed when the time series used to infer the covariance matrix become too short relative to the size of the portfolio [127–129], and the impact of having noisy covariance matrices in general [130–132]. The effect of having further constraints in the optimisation problem has also been explored, with very rich results. In particular, imposing that investors must provide a deposit proportional to the value of the underlying assets was shown to result in an exponential number of valid locally optimal portfolios, reminiscent of the number of metastable states in a spin-glass [42]. As argued in that paper, the existence of a very large number of nearly degenerate (or satisficing) solutions for a given optimisation problem is conceptually important because common knowledge can no longer be assumed. The complexity of the problems faced by the agents generates irreducible uncertainty, a quandary called *radical complexity*, and two strictly rational agents may then take very different decisions if their information and its treatment is not exactly the same. In that sense, radical complexity appears as a clear contradiction of the classical “homo economicus” and rational expectations paradigm discussed in Chap. 1.

In practice, a constraint that has long attracted much interest in the risk management industry consists in enforcing *long-only* portfolios. Denoting $\mathbf{w} \in \mathbb{R}^N$ the vector of weights associated to each of the N possible assets an investor can consider with $\sum_i w_i = 1$ (fully invested), a *long* position – which amounts to betting on the increase of the asset price – corresponds to a weight $w_i > 0$, and conversely a *short* sell – in which case we bet on the decrease of the price – refers to $w_i < 0$.¹⁹ There are a variety of reasons why one might need to avoid short sells, ranging from explicit investment mandates to extreme cases such as regulatory bans as

¹⁹While it might appear unusual to have negative weights in the portfolio, these can be understood rather intuitively. Indeed, to perform a short sell, one actually *borrow*s a stock to another participant, and sells it at the current price, in hope of buying it at a later time for a smaller price and returning it at that stage. The negative weight is then rather natural, as it indicates that one owes this fraction of their portfolio of stock i to the other participant.

those seen in Europe during the Coronavirus outbreak. On the fundamental level, very interesting behavior has been observed when the long-only constraint is enforced, as portfolios quickly tend to become very sparse, resembling choices made by individual stock pickers [133–135]. Interestingly, this constraint also allows to draw analogies with the ecological equilibria reached by some population dynamics models, which have also attracted considerable attention from the spin-glass community in recent years.

3.1.1 Problem statement

We consider a portfolio of N single assets, and assume (in line with the majority of studies on portfolio optimisation) that asset returns are correlated Gaussian variables. The portfolio statistics is fully characterized by its covariance matrix $C_{ij} = \langle \eta_i \eta_j \rangle - \mu_i \mu_j$, where we have introduced for each asset i the fluctuating return η_i and its expectation μ_i , and where this expectation is typically approximated as a time average in practice. The full correlation matrix is notoriously difficult to infer from noisy financial time series (see e.g. [132]), which is why simplifying hypotheses are generally used in the asset management literature. One of them is given by the one-factor risk model, which, rather than attempting to incorporate all possible sources of fluctuations, assumes that correlations are mostly due to the market exposure. Indeed, empirical analyses on stocks show that the top eigenvalue of the correlation matrix, corresponding to the so-called *market mode*, is largely dominant relative to the other eigenvalues (see e.g. [136]). Such an observation then suggests that all the assets have a common source of volatility (the overall market) explaining a large part of the cross-correlations, and leading to a covariance matrix of the form

$$C_{ij} = z_i \delta_{ij} + \beta_i \beta_j. \quad (3.1)$$

The asset-specific parameter $z_i = \sigma_i^2$, with σ_i the “idiosyncratic volatility”, models the uncorrelated asset volatility not accounted for by market-wide events, while β_i is to be understood as the sensitivity to the market (Chap. 9 of [17] for a complete discussion on the matter).²⁰

For an investor interested in constructing an optimal portfolio, the expected returns are of course key parameters. However for this theoretical analysis, which aims at drawing qualitative insights regarding the multiplicity of solutions and its implication on portfolio stability, we impose the simplification $\boldsymbol{\mu} = \mathbf{1}$, where $\boldsymbol{\mu} = \{\mu_i\}_{i \in [1, N]}$ (see Appendix A.6 for extensions to arbitrary μ_i ’s). In this case,

²⁰Unfortunately the correlation to the market is widely known as the “beta coefficient” β , which might induce some confusion with the intensity of choice or inverse temperature β . There is thankfully no such rationality parameter in this chapter, but we have adapted the typesetting to avoid confusion.

the optimal portfolio, also coined the Markowitz portfolio, which minimizes the risk $\sigma_p^2 = \sum_{i,j} C_{ij} w_i w_j$ for the given expected return $\mu_p = \sum_i w_i = 1$ (which conveniently also corresponds to the fully-invested constraint) enforced using a Lagrange multiplier, is given by:

$$w_i = \frac{\sum_j C_{ij}^{-1}}{\sum_{i,j} C_{ij}^{-1}}. \quad (3.2)$$

Within the one factor model, the covariance matrix can be easily inverted using the Woodbury matrix identity [137]. Up to a normalizing constant, one obtains:

$$w_i \propto \frac{1}{z_i} \left(1 - \beta_i \frac{\sum_j \beta_j / z_j}{1 + \sum_j \beta_j^2 / z_j} \right). \quad (3.3)$$

This equation is central for the problem that we aim to explore in the following.

3.1.2 Link with population dynamics

Let us take a minor sidestep towards ecological equilibria. Consider now $i = 1, \dots, N$ species associated to a certain “carrying capacity” in the environment. These species furthermore interact with each other, either competing for resources, or in predator-prey relationships, or else in a mutualistic, cooperative mode.

In its simplest form, where we consider that all species have identical growth rates $\mu_i = 1$, the population dynamics are described by the general Lotka-Volterra equations:

$$\dot{x}_i(t) = x_i(t) \left[\mu_i - \mu_i k_i^{-1} x_i(t) - \sum_j \alpha_{ij} x_j(t) \right], \quad \text{with } \mu_i = 1 \quad \forall i, \quad (3.4)$$

where x_i is the population of species i , k_i its carrying capacity, and α_{ij} is the $N \times N$ interaction matrix [41]. In this model, a positive entry α_{ij} corresponds to species i and j competing for resources or i being a prey and j being a predator (in which case $\alpha_{ji} < 0$). Setting $\dot{x}_i = 0$ to identify fixed points of the system yields the equilibrium population of the species:

$$x_i = \sum_j C_{ij}^{-1}, \quad (3.5)$$

where here $C_{ij} = z_i \delta_{ij} + \alpha_{ij}$, with $z_i = k_i^{-1}$. Naturally, in this context, one must have $x_i \geq 0 \quad \forall i$ since populations cannot be negative.

Experimentally, it is very difficult to gain insights on the nature of the interaction matrix or its eigenvalues. As a matter of fact, it is this observation that

initially motivated Robert May to use Random Matrix Theory arguments in his seminal paper [138]. While the qualitative phase portrait for the dynamical behavior of the model is independent of the exact distribution of α_{ij} [41], it seems natural from Eq. (3.5) that the equilibrium picture would be dependent on the interaction matrix model.

Here, we propose a drastic simplification and choose the interaction matrix to be of unit rank: $\alpha_{ij} = \beta_i \beta_j$, corresponding to species competing for a single common resource, in addition to the self-regulation included in Eq. (3.4). We will take β 's to be independent and identically distributed, with a probability density function $\rho(\beta)$. Consistent with the common notation in disordered systems, averages over this distribution (of disorder) will be indicated by an overline henceforth. The β_i coefficients then quantify how strongly species i competes for the unique resource with other species, and the interaction between two species then only depends on how strongly they both depend on the resource. With this model of interactions, the equilibrium populations map to the long-only optimal portfolio weights (up to a constant that does not affect the sign), and both problems can be treated identically based on Eq. (3.3). Note that the case of heterogeneous growth rates μ_i is equivalent to different average returns for stocks, and is discussed in Appendix A.6.

Taking $\bar{\beta} > 0$ is a natural choice, both to avoid placing ourselves in an unbounded growth regime, and more generally because ecosystems tend to be highly competitive. Naturally, the variance $\sigma^2 = \overline{\beta^2} - \bar{\beta}^2$ shall also play a key role in the equilibrium picture of the system and its properties. Finally, it is important to note that the unit rank model yields a symmetric interaction matrix, which amounts to either cooperation or pure resource competition. Predator-prey relations require, as noted above, asymmetric interactions between species i and j .

3.1.3 Spin-glass formulation

Suppose we now enforce the non-negativity constraint common to the two problems. For the portfolio, this means the positions (assets) associated to short sells after the Markowitz optimisation, i.e. associated to $w_i < 0$, will have to be removed from the portfolio altogether, reducing the effective universe from which stocks may be picked. Likewise, an extinct species ($x_i < 0$) is by definition removed from the ecological universe, leading to an ecology with a reduced number of viable species.

Here, we introduce “spins” $\{\theta\}$ that can take the value $\theta_i = 1$ if position i is included in the (possibly reduced) asset or species universe and $\theta_i = 0$ if it is excluded from it. Clearly, without excluding any specific entity, 2^N combinations

of $\{\theta\}$ can be constructed from the N assets or species initially considered²¹. The central question can therefore be reformulated as follows: we seek the number $\mathcal{N}_s$ of possible configurations $\{\theta\}$, among these 2^N , that satisfy the non-negativity condition (see Eq. (3.6) below). These spin variables are then related to the weights of the underlying positions through Eq. (3.3). Indeed, only the included positions, i.e. those with $\theta_i = 1$, now contribute to the sums, while the weights associate to positions with $\theta_i = 0$ are, by definition, discarded.

This quantity can easily be understood for financial assets, as it corresponds to the number of long-only Markowitz-optimal portfolios that can be constructed from a set of N assets. In the context of ecological equilibria, the interpretation is similar even if species are not “selected” in the same way as stocks. $\mathcal{N}_s$ can then be seen as the number of viable stable ecosystems that result from particular subsets of the N interacting species. The existence of solutions with a lower number of highly concentrated species in addition to the default, most diverse, solution is actually particularly interesting in the ecology context when considering the so-called *Allee effect*, which states that an increase in population density is correlated with higher survival probability [139]. In both cases, this quantity, which may appear somewhat artificial at this stage, will be essential in understanding how disorder chaos arises and can impact these systems in a very concrete way.

Naively, one could try to characterize the number of solutions by its average,

$$\overline{\mathcal{N}_s} = \overline{\sum_{\{\theta\}} \prod_k \Theta(\theta_k w_k)}, \quad (3.6)$$

where we take the convention $\Theta(0) = 1$ for the Heaviside step function, and averages are taken over the distribution of β . At this stage, one may notice how formally similar this enumeration is to the counting of TAP solutions discussed in the previous chapter, where the Heaviside step function replaces the Dirac δ distribution enforcing the TAP equations. As in the spin-glass setup, a key quantity in the study of the number of solutions is the annealed complexity

$$\Sigma = \frac{\log \overline{\mathcal{N}_s}}{N}. \quad (3.7)$$

As previously mentioned, in the Sherrington-Kirkpatrick spin-glass this quantity is indeed equivalent to its more representative quenched counterpart where the logarithm is averaged (requiring a more involved replica calculation) for metastable states of sufficiently high energy [106]. Another useful observable is the sparsity, describing the average fraction of the N possible entities that are included in the

²¹ $2^N - 1$ solutions to be exact, as the empty portfolio cannot satisfy the fully invested constraint.

configuration. We can write it as

$$u = \frac{1}{N} \sum_{i=1}^N \theta_i. \quad (3.8)$$

Among the $\mathcal{N}_s$ configurations satisfying the non-negativity constraint for a given set $\{\beta_i\}$, one will be “as full as it can be”, meaning u will reach its maximum value, the average value of which is noted m below, with $m \leq 1$. In asset management terms, this quantity corresponds to the most diversified long-only portfolios. For population dynamics, it is the most diverse ecosystem that can result from all possible viable ecosystems resulting from the N species.

3.1.4 Empirical data

In order to study the likely number of valid solutions, it is essential to have some information on the distribution of the β coefficients.

For the portfolio problem, where β is a widespread metric for an asset’s correlation relative to the market, high quality data is readily available. Using returns from a large number of US stocks over a two year time span (up to November 2020) reveals that the β s are normally distributed about 1, as shown in Fig. 3.1(a). We shall thus take as a starting point i.i.d. variables $\beta_i \sim \mathcal{N}(1, \sigma^2)$, which conveniently implies that σ and N are the problem’s sole parameters. Nevertheless, all calculations can easily be generalized to any $\bar{\beta} \neq 1$ since the problem is invariant under the simultaneous scaling of all β s and all z s by an arbitrary factor α and α^2 , respectively.

This being said, other distributions for β can also be of interest. In particular, if one focuses on specific sectors, the empirical distribution of volatility matches a uniform distribution relatively well as can be seen in Figure 3.1(b). Alternatively, looking at weekly returns rather than daily returns to construct the volatility gives thicker tails, between a Gaussian and Laplace distribution, as well as some slight skewness, visible in Fig. 3.1(c).

As mentioned in the previous section, there is unfortunately no such data for the interaction matrix in ecological communities, so in this context our model parameters should be considered with a grain of salt.

In the following numerical experiments, markers labeled by “Data” will be referring to calculations that are using the empirical distributions of β from Fig. 3.1. Practically, the shape and width of the empirical distributions are obtained by fitting the histograms, and data points are then constructed by random sampling with replacement for different values of N .




Figure 3.1: Distributions of β computed using the variance and covariance of returns averaged over one and two years respectively (as accurately estimating covariance requires more data). (a) Daily returns of the 1500 largest capitalisation stock, and fit to a normal distribution. (b) (dark green) Daily returns of the 120 largest capitalisation stocks in the energy and utilities industries, (light green) uniformly distributed points over the same interval for the same sample size. (c) Weekly returns of the 1000 largest capitalisation stocks, and fit to a generalized normal distribution of shape parameter $b \approx 3/2$.

3.2 Numerical experiments

We start by performing numerical experiments in order to find a relevant scaling regime. As we are interested by the static properties of a complicated optimization problem, there are no obvious dynamics to simulate in our problem. Instead, several approaches are possible.

3.2.1 Exact enumeration

As previously stated, our spin-glass inspired framing of the problem means that there are generically 2^N possible configurations to explore, giving as many unique Markowitz optimal solutions²² among which only a fraction are expected to satisfy the long-only constraint. The simplest approach is therefore to exhaustively explore these configurations, computing the Markowitz weights every time and determining whether there are any negative entries in the solution vector. This “bruteforce” approach, which will be referred to as *exact enumeration*, is obviously exponentially expensive computationally, and can therefore only be carried out for small to intermediate system sizes ($N = 32$ already gives over 4 billion Markowitz portfolios to compute!). However, it has the advantage of giving access to all properties of the solution space, allowing for a precise comparison of the properties of valid solutions.

²²In the following, a solution refers to any non-negative $\mathbf{w}$ for which Eq. (3.2) taken on the covariance matrix *restricted* to the indices where $\theta_i = 1$ is satisfied

Chapter 3. Rationality versus complexity: the example of portfolio optimization

First, one can consider the quadratic form

$$\mathcal{H}(\{\beta\}, \{\theta\}) = \sum_{i,j=1}^N C_{ij} \theta_i w_i \theta_j w_j, \quad (3.9)$$

with the weights w_i solutions to Eq. (3.3) with summations restricted to positions with $\theta_i = 1$, represents the metric that is minimized under constraint, i.e. the energy of a configuration in physical terms. For the portfolio problem, this is obviously the total portfolio square volatility. Although less straightforward, $\mathcal{H}$ can also be understood intuitively in the population dynamics context. Indeed, as C_{ij} represents the level of competition for the resource between species i and j , the quantity to be minimized corresponds to the aggregated level of competition (including self-competition) for the surviving species.

The second quantity of interest is the overlap between two configurations $\{\theta\}$ and $\{\tilde{\theta}\}$. We choose to define it as

$$q(\{\theta\}, \{\tilde{\theta}\}) = \frac{1}{N} \sum_{i=1}^N \theta_i \tilde{\theta}_i - \left(\frac{1}{N} \sum_{i=1}^N \theta_i \right) \left(\frac{1}{N} \sum_{i=1}^N \tilde{\theta}_i \right). \quad (3.10)$$

This definition differs from the usual spin-glass expression with the addition of the second term on the rhs, that is included to ensure that two statistically independent configurations have zero overlap on average.

Using these two metrics, we can study the distribution among valid configurations of their excess variance or level of competition relative to the value for the globally optimal configuration $\{\theta^*\} = \operatorname{argmin}_{\{\theta\}} \mathcal{H}(\{\beta\}, \{\theta\})$ for a given draw $\{\beta\}$,

$$\Delta\mathcal{H} = \mathcal{H}(\{\beta\}, \{\theta\}) - \mathcal{H}(\{\beta\}, \{\theta^*\}) \quad (3.11)$$

as well as their normalized overlap with this global optimum

$$\mathcal{Q} = \frac{q(\{\theta\}, \{\theta^*\})}{q(\{\theta^*\}, \{\theta^*\})}. \quad (3.12)$$

A value of $\mathcal{Q}$ close to 1 indicates a configuration has a large number of common species or assets with the optimal configuration, $\mathcal{Q}$ close to 0 corresponds to solutions much sparser than the optimum while negative values of the overlap are reached for configurations that are largely full but orthogonal to the best possible outcome.

The distributions of $\Delta\mathcal{H}$ and $\mathcal{Q}$ can be obtained directly from the exact enumerations for small values of N repeated over several draws of the disorder. Such a result for normally distributed β is shown in Fig. 3.2, which also displays the joint density of the two quantities. Looking at the overlap, we find that, as one




Figure 3.2: Solution space explored by exact enumeration for Gaussian β with $N = 28$, $\sigma = 0.1$, averaged over 30 realizations. Top left: distribution of overlap relative to the optimal configuration Q . Bottom left: distribution of excess variance relative to the optimal configuration $\Delta\mathcal{H}$. Right: heatmap of the joint density of these two quantities.

could have expected, the majority of valid solutions are composed of a relatively small number of non-zero spins and therefore have Q close to zero. More surprisingly, looking at $\Delta\mathcal{H}$ reveals that a very large fraction of these are associated to a small excess variance or level of competition relative to the minimum. This is further confirmed by the joint distribution, where we indeed observe that many configurations display a small value of $\Delta\mathcal{H}$ despite having $Q \approx 0$. As such, the multiplicity of solutions and associated complexity is of a great importance in this problem. Not only do we find a large number of portfolios or ecosystems that satisfy the constraint, a large fraction of these achieve a portfolio variance or level of competition very close to the best possible outcome. These quasi-degenerate solutions might therefore become optimal following a small change in the disorder. This idea is at the root of the disorder chaos investigated in Section 3.5, and related to the *de facto* limitation of rational choice arguments in complex situations.

3.2.2 “Pruning” algorithm

As previously stated, exactly enumerating the long-only solutions for a given draw of disorder is a powerful approach, as it allows the precise study of all valid solutions, but comes at a very high computational cost. To explore larger N that are relevant for practical applications and the impact of the system size on the optimal solution in an analytically tractable regime, we turn to a greedy heuristic which turns out to be exact in this case.

The algorithm simply consists in removing iteratively entries for which Eq. (3.3) gives negative weights, i.e. setting $\theta_i = 0$ for these positions, until all posi-

tions are acceptable (see [140]). In the population dynamics context and given we are considering homogeneous growth rates, this “pruning” approach is in fact a very straightforward Darwinian evolution in the sense that most vulnerable species (i.e. those with the largest β_i) are extinct first, until all species in the ecosystem survive (and represent a positive fraction of the population). The outcome of the algorithm is then the most diverse solution possible, for which the fraction of non-zero entries is the previously defined sparsity m . The pruning algorithm will therefore allow us to probe the evolution of this quantity for different values of σ and N for much larger systems than what was possible with an exhaustive search.

The result for stock-compatible $\beta \sim \mathcal{N}(1, \sigma^2)$ in Fig. 3.3(a) is consistent with the findings of Lehalle and Simon [134]: the sparsity decreases rapidly and non-trivially when the β s cease to be very tightly distributed about unity. Interestingly, the sparsity clearly appears to be a function of the parameter $\chi = \sigma N$ only (see Fig. 3.3(a)). Such a scaling provides precious insight for the analytical formulation of the problem, as will become apparent in the following section. It should be noted that such a scaling result ceases to hold for large values of σ , as a large standard deviation yields a significant fraction of assets with negative β that can obviously be included in the long-only portfolios (see Eq. (3.3)), thereby causing m to increase again at large σ .

This effect can be partially observed in the empirical points that have a slightly wider distribution of β ($\sigma \approx 0.3$) and a few negative entries (not shown). This being said, in a range for N and σ relevant for applications, the evolution of m for real data points appears to be roughly in line with the $\chi = \sigma N$ scaling curve where the fully numerical points lie.

The same procedure may be repeated for uniformly distributed β as show in Fig. 3.4(a). The result is qualitatively very similar, albeit with a slower decrease in m with N . Interestingly we recover the scaling $\chi = \sigma N$ where σ now governs the width of the distribution. Like in the Gaussian case, the points sampled from empirical data appear to be slightly too widely distributed to perfectly match the points from continuous probability densities, although the evolution of m appears to be very close up to some offset.

Note that the pruning algorithm can be slightly modified in order to explore a larger region of the solution space, beyond the globally optimal solution. Indeed, by introducing some stochasticity in the iterative process, i.e. by setting $\theta_i = \{0, 1\}$ for stocks with negative and non-negative weights respectively *with probability*, the algorithm will evolve towards some slightly sub-optimal (and therefore less diverse) yet still valid solutions. Depending on the probability with which *a priori* valid and non-valid stocks or species are excluded and included in the solution, one will explore solutions with a more or less large overlap with the best possible outcome. In other words, playing with this probability will shed light on different regions of the joint distribution presented in Fig. 3.2.




Figure 3.3: Numerical and theoretical results for $\beta \sim \mathcal{N}(1, \sigma^2)$. (a) Maximum sparsity m as a function of $\chi = \sigma N$, inset focusing on the large χ region. (b) Number of solutions $\langle N_s \rangle$ as a function of N , obtained by exact enumerations averaged over relatively few (4-10) samples. Straight lines display the respective theoretical predictions, dotted line corresponds to 2^N . (c) Complexity Σ as a function of σN resulting from the exact enumerations and normalized by $\log 2$, inset zooming out to show the large χ region. The full line is the numerically exact result, and the dotted line is an asymptotic approximation based on Eq. (3.27) below. The legends are shared for (b) and (c).




Figure 3.4: Numerical and theoretical results for $\beta \sim \mathcal{U}(1 \pm \sqrt{2}\sigma)$. See Fig. 3.3 for the detailed caption.

3.3 Analytical setup

3.3.1 Self-consistent equation

While it may appear natural to attempt to directly tackle the enumeration of valid solutions using Eq. (3.6) and a Fourier representation of the Heaviside function, such a calculation quickly requires a Gaussian assumption on the distribution of β and thus lacks generality. The alternative taken here is to first study directly the maximum sparsity m as a function of N , before translating this quantity back to the number of long-only portfolios.

Going back to Eq. (3.3) that relates the weight of position i to its β_i , it immediately appears that there should be a threshold value β^+ above which a position will likely be shorted (or a species go extinct), which must thus be excluded. This is quite reasonable intuitively: an investor wishing to take the least possible risk and unable to balance volatility through shorts will be unlikely to pick excessively risky stocks or bonds. Likewise, since large values of β are associated to species subject to increased competition, those species are likely to go extinct and thus not be present in the equilibrium population. Given the distribution $\rho(\beta)$, the average maximum sparsity m is then related to the threshold β^+ through:

$$m = \int_{-\infty}^{\beta^+} d\beta \rho(\beta). \quad (3.13)$$

Therefore, calculating β^+ will directly yield m . Of course, this threshold is only valid in a statistical sense, and for a given set of β 's its value will differ from the mean. Writing $\tilde{\beta}^+$ the fluctuating variable representing the largest β to be included for a unique realization of the disorder, and modifying Eq. (3.3) with the previously introduced spin notation directly gives

$$\tilde{\beta}^+ = \frac{\sum_j \beta_j^2 \theta_j / z_j + 1}{\sum_j \beta_j \theta_j / z_j}, \quad (3.14)$$

where we now have $\theta_i = 1$ for $\beta_i < \tilde{\beta}^+$ and $\theta_i = 0$ otherwise. In order to recover the typical behavior of interest, we introduce the probability of inclusion conditioned to β . We define $\text{Prob}(\theta_i = 1 | \beta_i) = F_N(\beta_i)$, where $F_N(\beta)$ is a smooth step-like function, centered about the mean β^+ . Clearly, we require $F_N(\beta)$ to be monotonously decreasing, with

$$F_N(-\infty) = 1 \quad \text{and} \quad F_N(\infty) = 0, \quad (3.15)$$

and $F'_N(\beta)$ is therefore peaked in a region around β^+ , the width of which is expected to decrease as N increases.

Chapter 3. Rationality versus complexity: the example of portfolio optimization

At this stage, we start by assuming that two spins are uncorrelated at order N^{-1} , which shall be checked *a posteriori* by ensuring the Onsager “reaction” term is $o(N^{-1})$. The sums at the numerator and denominator of Eq. (3.14) may therefore be treated using the central limit theorem, i.e. for $N \gg 1$

$$\frac{1}{N} \sum_j \frac{\beta_j^k \theta_j}{z_j} \approx \langle z^{-1} \rangle \langle \beta^k \rangle + \frac{1}{\sqrt{N}} \xi_k \quad (3.16)$$

where the partial expectation operator $\langle \langle \dots \rangle \rangle$ is defined, for an arbitrary test function $g(\cdot)$, as

$$\langle \langle g(\beta) \rangle \rangle := \int_{-\infty}^{\infty} d\beta g(\beta) \rho(\beta) F_N(\beta) \quad (3.17)$$

and ξ_k is a zero-mean Gaussian variable, with a variance that depends on k . Eq. (3.14) may then be rewritten as

$$\tilde{\beta}^+ = \beta^+ + \frac{1}{\sqrt{N}} \xi, \quad (3.18)$$

where the full expressions of β^+ and ξ are given in Appendix A.1. From the very definition of $F_N(\beta)$, which represents the probability of inclusion in the reduced universe of an asset or species with respectively correlation or interaction strength β , one can write $F_N(\beta) = \text{Prob}(\beta < \tilde{\beta}^+)$. Using that ξ is a Gaussian noise, this can be expressed as

$$F_N(\beta) = \frac{1}{2} \text{erfc} \left[\frac{\sqrt{N}(\beta - \beta^+)}{\gamma\sqrt{2}} \right], \quad (3.19)$$

where γ is the standard deviation of ξ .

We now place ourselves in the scaling regime where $\sigma = \chi/N$ (motivated by numerical results). In this case the width of the distribution of $\beta - 1$ scales as N^{-1} . Assuming that in this regime $\gamma \rightarrow 0$, one finds that to leading order $\langle \langle \beta^k \rangle \rangle = m + O(N^{-1})$. Now, given the expression for γ^2 in App. A.1, one finally obtains

$$\gamma = O(N^{-1/2}), \quad \text{when } \sigma = O(N^{-1}), \quad (3.20)$$

which justifies our assumption that $\gamma \rightarrow 0$ for large N . It furthermore shows that the width of the smoothed step function $F_N(\beta)$ (Eq. (3.19)) scales as N^{-1} .

This result then allows us to explicitly make Sommerfeld-like expansions of averages, as described in App. A.2, that now have no contribution at order N^{-1} . Eliminating the higher order terms appropriately finally yields the equation for the mean threshold

$$\beta^+ = \frac{\langle \beta^2 \rangle_c}{\langle \beta \rangle_c} + \frac{1}{N} \frac{\bar{z}}{\langle \beta \rangle_c} + O\left(\frac{1}{N^2}\right), \quad (3.21)$$

valid in the regime of interest $\sigma = \chi/N$, with $\bar{z} = \langle z^{-1} \rangle^{-1}$ and

$$\langle g(\beta) \rangle_c := \int_{-\infty}^{\beta^+} d\beta g(\beta) \rho(\beta). \quad (3.22)$$

Eq. (3.21) is self-consistent in the sense that β^+ appears in both sides of the equation.

Recall that this equation for β^+ assumes negligible correlations between the occupation variables θ_i . In the spirit of a *bona fide* cavity calculation, one should look at the effect of the introduction of an additional asset or species on the already existing θ_i . Knowing the importance of the Onsager reaction term in Sherrington-Kirkpatrick spin-glasses, that turns the naive mean-field equation into the celebrated TAP equation (Sec. 2.2.2), it is important to ensure that the average threshold is not affected by a similar term. This is done in App. A.3 where we check that the introduction of a new spin does not alter the above equation at order N^{-1} . As such, Eq. (3.21) is our central analytic result for the problem at hand, which we shall solve for different distributions $\rho(\beta)$ in Sec. 3.4.

3.3.2 Complexity and number of solutions

Now, we define $\mathcal{N}(K, N)$ to be the average number of solutions satisfying the constraint with K among the N possible spins included. We may write an iterative equation to describe the evolution of this quantity as $N \rightarrow N + 1$. First, the addition of this new element – that we will take to be at index 0 and associated to β_0 – is only possible if β_0 is small enough. If we recall the probabilistic interpretation of the maximum sparsity $m = \text{Prob}(\beta_0 \leq \beta^+)$, the probability of $\theta_0 = 1$ being compatible with the constraint is simply given by $m(\sigma, K)$. In order to form such a solution, with K among the now $N + 1$ spins included, the new element must be added to a solution previously comprising $K - 1$ spins. However, a fraction of the solutions with $K - 1$ nonzero spins are rendered invalid due to the fact that β^+ is a decreasing function of K . Those positions are such that $\beta^+(K) < \beta_i < \beta^+(K - 1)$, and occur with probability

$$p(\sigma, K) = \int_{\beta^+(K)}^{\beta^+(K-1)} d\beta \rho(\beta) = m(\sigma, K - 1) - m(\sigma, K), \quad (3.23)$$

and given β s are drawn independently, we finally find the expression

$$\mathcal{N}(K, N + 1) = \mathcal{N}(K, N) + m(\sigma, K)[1 - p(\sigma, K)]^{K-1} \mathcal{N}(K - 1, N) \quad (3.24)$$

to describe the evolution of the number of solutions with K non-zero spins as N increases. To properly initialise and close the recursion, we require

$$\mathcal{N}(0, 0) = 1 \quad \text{and} \quad \mathcal{N}(N + 1, N) = 0. \quad (3.25)$$

Chapter 3. Rationality versus complexity: the example of portfolio optimization

The quantity that interests us, the average total number of solutions satisfying the constraint, is then simply given by

$$\overline{\mathcal{N}}_s = \sum_{K=1}^N \mathcal{N}(K, N). \quad (3.26)$$

Defining $n(x, t)$ to be the continuous analogue of $\mathcal{N}(K, N)$ with $K \rightarrow x$ and $N \rightarrow t$, the iterative equation may be rewritten as a partial differential equation, valid in the large N limit. To leading order, i.e. neglecting a diffusion term of order N^{-1} , one has

$$\partial_t n(x, t) + e^{x\varphi'(\sigma x)} \varphi(\sigma x) \partial_x n(x, t) = e^{x\varphi'(\sigma x)} \varphi(\sigma x) n(x, t) \quad (3.27)$$

where we have used the scaling result $m(N, \sigma) = \varphi(\chi)$ with $\chi = \sigma N$ as observed in numerical experiments, and further justified by the analytical calculations in the next section.

This inhomogeneous advection equation may then be treated with the method of characteristics [141]. Taking the characteristic curve s in (x, t) space, and writing $z(s) = n(x(s), t(s))$ the solution along the curve, the problem reduces to the system of ordinary differential equations

$$\frac{dt}{ds} = 1 \quad (3.28)$$

$$\frac{dx}{ds} = e^{x(s)\varphi'(\sigma x(s))} \varphi(\sigma x(s)) \quad (3.29)$$

$$\frac{dz}{ds} = e^{x(s)\varphi'(\sigma x(s))} \varphi(\sigma x(s)) z(x(s), t(s)), \quad (3.30)$$

with boundary conditions

$$t(0) = 0, \quad x(0) = 0, \quad z(0) = 1. \quad (3.31)$$

The solution satisfying these boundary conditions then directly corresponds, for $t = N$, to the dominating term in the sum given in Eq. (3.26).

To summarize, the self-consistent equation (3.21) allows one to determine the average threshold β^+ for the inclusion of an asset or species in the non-negative solution for a given sparsity. This quantity will in turn yield the expression of the maximum sparsity $m(N, \sigma) = \varphi(\chi)$ in the regime $\sigma = \chi/N$. Solving the set of characteristic equations tracing back to the known boundary conditions shall finally give an expression of the average number of solutions, and therefore the annealed complexity.

3.4 Distribution-specific results

3.4.1 Gaussian disorder

As argued with the data presented in Fig. 3.1(a), taking β to be normally distributed with mean 1 and variance σ^2 appears to be a good approximation for the portfolio problem. Going back to Eq. (3.21), all the terms of interest can be written exactly using the Gaussian cumulative distribution function $\Phi(x) = \frac{1}{2}(1 + \operatorname{erf}(x/\sqrt{2}))$. Taking $\sigma = \chi/N$ and introducing the *ansatz*

$$\beta^+ = 1 + \frac{\chi f(\chi)}{N} \quad (3.32)$$

allows one to rewrite the self-consistent equation as

$$\chi f(\chi) = \frac{\bar{z}}{m} - \frac{1}{m} \frac{\chi}{\sqrt{2\pi}} e^{-\frac{1}{2}f(\chi)^2}, \quad (3.33)$$

with $m = \varphi(\chi) = \Phi(f(\chi))$. As anticipated in the previous section, $m(N, \sigma)$ indeed only depends on χ in the scaling regime.

Setting $\bar{z} = 1$ (without loss of generality, since it simply corresponds to the rescaling $\chi \rightarrow \chi/\bar{z}$) this equation may be solved numerically for f at given χ , the result of which is plugged back into the expression for m and is shown by the continuous line in Fig. 3.3(a). This analytical result is in excellent agreement with numerical experiments, which gives us confidence that our self-consistent equation is exact in the regime of interest. As expected, while qualitatively reasonable, the model does not perfectly describe the sparsity corresponding to more broadly distributed empirical β 's.

This numerically obtained $m = \varphi(\chi)$ can also be injected in the iterative expression for $\mathcal{N}(K, N)$ given in Eq. (3.24). Summing all contributions, the mean number of solutions $\overline{\mathcal{N}}_s$ and associated complexity Σ are computed and shown by the continuous lines in Fig. 3.3(b-c). The match between this semi-analytical solution and the numerical results is also excellent, this time for both the arbitrary and empirically determined values of σ .

Based on the numerical solution of Eq. (3.33), we find that $f(\chi \gg 1)$ quickly reaches large negative values. The error functions through which φ is expressed can therefore be approximated asymptotically through the method of steepest descent. Keeping the first two terms in the series expansion of $m = \varphi$ in the self-consistent equation, and taking iterated logarithms, one finally finds, at the leading order in the scaling regime:

$$\varphi(\chi) \approx \frac{\sqrt{2 \log \chi}}{\chi}, \quad (\chi \gg 1) \quad (3.34)$$

Chapter 3. Rationality versus complexity: the example of portfolio optimization

This asymptotic result is compared to the numerical experiments in the inset of Fig. 3.3(a), displaying a very good fit for values as small as $\chi \sim 10$.

As detailed in App. A.4, this result may be used in the characteristic Eq. (3.29). At the leading order, we find the expression of x along the characteristic

$$x(s) = \sqrt{\frac{2s}{\sigma}} (\log \sigma s)^{\frac{1}{4}} \left[1 + O\left(\frac{\log \log \sigma s}{\log \sigma s}\right) \right]. \quad (3.35)$$

Eq. (3.30) may then be integrated to find $z(s)$ the number of solutions along the characteristic,

$$\log z(s) = \sqrt{\frac{2s}{\sigma}} (\log \sigma s)^{\frac{1}{4}} \left[1 + O\left(\frac{1}{\log \sigma s}\right) \right] \quad (3.36)$$

From Eq. (3.28) and the associated boundary condition, we may now finally set $s = t = N$. Going back to original variables of the problem, we therefore have the asymptotic evolution for the number of non-negative solutions

$$\mathcal{N}_s \sim \exp \left(\sqrt{\frac{2N}{\sigma}} (\log \sigma N)^{\frac{1}{4}} \right) \quad (3.37)$$

and the associated annealed complexity

$$\Sigma \approx \sqrt{2} \frac{(\log \chi)^{\frac{1}{4}}}{\sqrt{\chi}}. \quad (3.38)$$

This fully analytical asymptote is compared to the previously obtained numerically exact solution of the recursion relation (dubbed “semi-analytical” below) in the inset of Fig. 3.3(c). The result appears satisfactory, although the conclusions are limited by the numerical difficulty of obtaining the semi-analytical result for large values of N . A careful observation suggests a small shift between the two curves, which might be explained by the second derivative (diffusion) term in the partial differential Eq. (3.27), which we neglected.

In any case, this result corresponds to a growth slightly faster than $e^{\sqrt{N}}$ but significantly slower than e^N : asymptotically, the complexity of the rank-one portfolio problem, or of the rank-one ecological problem, is zero, contrarily to the spin-glass case. But the average number of different possible solutions is still very large when N is large.

It immediately appears however, that this solution is somewhat contradictory with the previously found behavior of the maximum sparsity. Indeed, taking a closer look at the solution for $x(s)$ along the characteristic curve, we find that the associated sparsity $u^* = x(t)/t$ is given by

$$u^* \sim \frac{(\log \chi)^{\frac{1}{4}}}{\sqrt{\chi}}. \quad (3.39)$$

Comparing with Eq. (3.34), we find that $u^* \gg \varphi(\chi)$ for $\chi \gg 1$.

In other words, it appears that the configurations dominating the count of the mean number of solutions are those with a number of non-zero spins K^* greatly exceeding the theoretical prediction mN . While surprising at first, this result means that the average (over β s) number of solutions is dominated by extremely rare configurations $\{\beta_i\}$ that allow $K^* \gg mN$ positions to survive in the portfolio. Even if rare, such configurations allow an exponentially large number of portfolios to exist, i.e. $C_N^{K^*}$. Hence the distribution of $\mathcal{N}_s$ is heavily skewed towards large values, corresponding to events that are extremely unlikely to be witnessed in reality. For typical configurations of the β_i 's, on the other hand, one expects that eligible portfolios are much smaller, and contain at most mN assets. Correspondingly, the typical number of solutions is expected to be of order of $e^{Nm} \ll \overline{\mathcal{N}_s}$. In order to compute precisely the typical number of solutions, one should compute $\log \overline{\mathcal{N}_s}$ and the associated quenched complexity. This would require going back to the direct formulation given in Eq. (3.6), expressing the Heaviside step function with its Fourier representation and making use of the replica trick as detailed in numerous works relating to spin-glasses (see [44, 106]). We leave such a calculation for later investigations. The computation of the probability of observing these rare configurations that dominate the average solution count, that is likely directly dependent on the distribution of the β_i 's, is also left for future work.

This difference between the typical (and average) maximum sparsity m and the most likely effective sparsity u^* resulting from the iterative procedure is apparent in Fig. 3.5 that considers normally distributed β , and is comparable in the uniform case. Here, we stress that the former corresponds to the self-averaging fraction of positions that may be included in a long-only Markowitz optimal portfolio, while the latter corresponds to the fraction of occupied positions associated to the portfolio that dominated the calculation of $\overline{\mathcal{N}_s}$ at large N (i.e. the solution at the saddle in the spin-glass language). For small χ , we have $u^* = \frac{1}{2} < m$, as naively expected. However, beyond $\chi \approx 5$ we find that indeed the mean behavior u^* exceeds the typical sparsity m . In this region, we would therefore expect the calculation of $\overline{\mathcal{N}_s}$ and the associated annealed complexity Σ to deviate from that observed in a moderate number of numerical experiments. The divergence between u^* and m appears to be relatively slow however, explaining why it is not clearly noticeable in the numerical results in Fig. 3.3 and 3.4, where only a small fraction of elements are excluded as m is still close to unity.

3.4.2 Uniform disorder

We now extend the results to other distributions of β . The uniform distribution is an interesting case, as it can model the case of stocks restricted to certain industries (see 3.1(b)).




Figure 3.5: Evolution of the effective sparsity of the maximum of $\mathcal{N}(K, N)$ calculated iteratively for Gaussian β of variance σ^2 , up to $N = 400$ (points), compared to the maximum sparsity (line). The red vertical line separates a small χ regime where one expects that $\bar{\mathcal{N}}_s \approx \mathcal{N}_s^{\text{typ}}$, from a large χ regime where $\bar{\mathcal{N}}_s \gg \mathcal{N}_s^{\text{typ}}$ when $N \rightarrow \infty$.

In the uniform case, the distribution is still centered about $\beta = 1$ but now has width $2\sqrt{2}\sigma$. Once again, we take the scaling $\chi = \sigma N$ from numerical experiments, and the *ansatz* $\beta^+ = 1 + \chi f(\chi)/N$. The moments up to the threshold β^+ assuming $\beta^+ < 1 + \sqrt{2}\sigma$ may then be easily written explicitly given the simple expression of the uniform distribution. Taking the self-consistent equation at order N^{-1} , where we can once again take $\bar{z} = 1$ without loss of generality, finally gives the expression for the function $f(\chi) = -\sqrt{2} \pm 2^{5/4}/\sqrt{\chi}$, from which the maximum sparsity directly follows by picking the solution giving the positive result. This solution requires the threshold to be before the right edge of the distribution, and hence must be completed with the result beyond which saturates the maximum value $m = 1$. Combining both gives a closed form solution for the entirety of the domain without having to rely on asymptotics

$$\varphi(\chi) = \begin{cases} 1 & \text{for } \chi \leq \frac{1}{\sqrt{2}}, \\ \frac{1}{2^{1/4}\sqrt{\chi}} & \text{for } \chi > \frac{1}{\sqrt{2}}. \end{cases} \quad (3.40)$$

This solution corresponds to the continuous line in Fig. 3.4(a). Once again, the match with numerical simulations is very good, whereas – as discussed in Section 3.2 – there is a small offset relative to the empirically sampled points that lie slightly outside of the analytically tractable region. Note that the typical sparsity of the portfolios decreases with N much more slowly in the uniform case than in the Gaussian case.

The fully analytical solution for $m = \varphi(\chi)$ is substituted in the iterative formula for the mean number of solutions, resulting in the continuous lines in Fig. 3.4(b-c). Clearly, this theoretical result displays a very good match with the numerical points across all values of σ tested.

Chapter 3. Rationality versus complexity: the example of portfolio optimization

As before, we now employ this expression in the set of ordinary differential equations to solve the partial differential equations describing the evolution of the number of solutions. Thanks to the simple expression for φ that is now valid for all values of χ , the integration may be carried out with no difficulty (A.4), giving

$$\log z(s) = \left(\frac{3}{2^{5/4}} \frac{s}{\sqrt{\sigma}} \right)^{\frac{2}{3}} \left[1 + O\left(\frac{1}{\sqrt[3]{\sigma s}} \right) \right] \quad (3.41)$$

and thus simply replacing $s = t = N$,

$$\overline{\mathcal{N}}_s \sim \exp \left\{ \left(\frac{3}{2^{5/4}} \frac{N}{\sqrt{\sigma}} \right)^{\frac{2}{3}} \right\}. \quad (3.42)$$

As we might have expected from the slower decrease in maximum sparsity relative to the Gaussian result, the average number of solutions grows faster in the uniform case. The annealed complexity is now asymptotically given by

$$\Sigma \sim \left(\frac{3}{2^{5/4}} \frac{1}{\sqrt{\chi}} \right)^{\frac{2}{3}}, \quad (3.43)$$

that is plotted with the dashed line in the inset of Fig. 3.4(c). Here, the fully analytical expression appears more or less in line with the semi-analytical iterative solution.

As for the Gaussian case, we notice that the sparsity of the configurations dominating the enumeration is given by $u^* = x(t)/t \sim \chi^{-\frac{1}{3}} \gg \varphi(\chi)$. Just as before, we have therefore calculated a mean number of solutions that appears to greatly exceed the typical result observed. The typical behavior would then also require to compute $\overline{\log \mathcal{N}}_s$, which in this uniform case would not be as similar to typical spin-glass calculations that rarely, if ever, involve uniform distributions with a finite support. We note however that the typical number of solutions in this case should grow as $\exp(\sqrt{N})$, i.e. much faster than in the Gaussian case where it only grows as $\exp(\sqrt{\log N})$.

3.4.3 Bridging the gap: generalized normal distribution

To understand why two different decays in maximum sparsity hold for the normal and uniform distributions, we use of the generalized normal distribution

$$\rho_b(\beta) = \frac{b}{2\sqrt{2}\sigma\Gamma(1/b)} e^{-\left(\frac{|\beta-1|}{\sigma\sqrt{2}}\right)^b} \quad (3.44)$$

where b is a shape parameter allowing to recover the usual normal distribution of unit mean and standard deviation σ^2 for $b = 2$, and the uniform distribution

Chapter 3. Rationality versus complexity: the example of portfolio optimization

centered at 1 and of width $2\sqrt{2}\sigma$ by taking the limit $b \rightarrow \infty$. Moreover, probing $b \leq 1$ can provide insights on the problem with heavier tailed distributions of β 's, $b = 1$ corresponding to the Laplace case, which may be of interest when considering e.g. the weekly returns presented in Fig. 3.1(c).

The first step in our search for an analytical solution in this general formulation is to express the moments up to the threshold $\langle \beta \rangle_c$ and $\langle \beta^2 \rangle_c$ as well as m itself in workable forms. As detailed in App. A.5, this can be done by reintroducing the expressions for σ and β^+ . The generalized self-consistent equation now reads

$$\chi f(\chi) = \frac{\bar{z}}{m} - \frac{\chi}{m} \int_{-f(\chi)}^{\infty} du u \rho(u), \quad (3.45)$$

where $u = (\beta - 1)/\sqrt{2}\sigma$ and we have postulated $f(\chi) < 0$ which is intuitive from the expression of β^+ (we expect the threshold to be smaller than the mean value of β , regardless of the distribution). Taking $b = 2$, the integral can be evaluated exactly and we recover Eq. (3.33) as expected. As for the two previous cases, setting $\bar{z} = 1$ simply corresponds to rescaling $\chi \rightarrow \chi/\bar{z}$.

For b sufficiently small, we may approximate the integrals asymptotically as we expect $f(\chi)$ to have a large magnitude for these widely distributed β 's. The resulting self-consistent equation (App. A.5) now reads

$$e^{(\frac{|f(\chi)|}{\sqrt{2}})^b} = \frac{\chi}{\sqrt{2}b\Gamma(1/b)} \left(\frac{|f(\chi)|}{\sqrt{2}} \right)^{2-2b} \quad (3.46)$$

while the maximum sparsity at the leading order is

$$m = \varphi(\chi) = \frac{e^{-(\frac{|f(\chi)|}{\sqrt{2}})^b}}{2\Gamma(1/b)} \left(\frac{|f(\chi)|}{\sqrt{2}} \right)^{1-b}. \quad (3.47)$$

Introducing the variables

$$y = \left(\frac{|f(\chi)|}{\sqrt{2}} \right)^b \quad \text{and} \quad x = \frac{\chi}{\sqrt{2}b\Gamma(1/b)}, \quad (3.48)$$

the self-consistent equation takes the much simpler form

$$y^{2-\frac{2}{b}} e^y = x, \quad (3.49)$$

giving in turn $2\Gamma(1/b)\varphi = y^{1-\frac{1}{b}}/x$. For a given value of b , this simplified self-consistent equation can be either solved semi-analytically or asymptotically in the limit of χ and therefore x large. For instance taking $b = 2$, Eq. (3.49) gives $y = W(x)$ the Lambert W function and thus

$$m = \frac{\sqrt{W(x)}}{2\sqrt{\pi}x} \sim \frac{\sqrt{2\log \chi}}{\chi}, \quad (3.50)$$




Figure 3.6: Scaled sparsity as a function of b , σ and N for generalized normal distributions. The continuous, dashed and dot-dashed lines correspond to the uniform, Gaussian and Laplace asymptotes respectively.

thereby recovering the previously obtained result. Interestingly, the case $b = 1$ corresponding to Laplace distributed β 's yields the exact relation

$$m = \frac{1}{2x} = \frac{1}{\sqrt{2}\chi}, \quad (3.51)$$

suggesting $\overline{\mathcal{N}}_s \sim e^{\sqrt{N}}$, slightly slower than for Gaussian β 's. Both asymptotic solutions are shown in Fig. 3.6, displaying a good match as χ increases.

For $b \gg 1$, the problem is not as straightforward. Indeed, as the shape of the distribution approaches the uniform case, the effective support narrows to reach sharp cutoffs at $\beta = 1 \pm \sqrt{2}$ for $b \rightarrow \infty$. In this limit, we therefore require $-f(\chi) < \sqrt{2}$, which is obviously unsuitable for the previously taken asymptotic approximations of the integrals. Instead, the limit $b \rightarrow \infty$ must be taken before evaluating the integral. Doing so, one can directly recover the uniform expressions for $f(\chi)$ and m from the previous section, and therefore $m = (\sqrt{2}\chi)^{-\frac{1}{2}}$.

Rescaling the $b \rightarrow \infty$ result suggests $2\varphi\Gamma(1/b) \approx \sqrt{2}/\sqrt{x}$ for large b . This large b solution, and the crossover between the two regimes, with φ decaying as $\chi^{-\frac{1}{2}}$ and χ^{-1} respectively, can be seen in Fig. 3.6. Interestingly, for all finite values of b the second regime will be reached eventually as χ is increased, and only the uniform distribution will remain in the first regime, the slower decay of which translates in a larger number of solutions. As such, the uniform distribution will be the case within the generalized normal family allowing for the largest number of solutions, with finite $b > 2$ cases only affecting the exponent of the logarithmic term in the complexity.

While the above theoretical setup should hold for finite $b < 1$, it is difficult to avoid a large number of negative β 's when considering thicker tails, at which point we would cease to observe a monotonous decrease of m in χ . Besides, it seems unlikely that a heavy tailed distribution of infinite support would correctly depict the distribution of asset correlations. For interacting species, negative interactions are not as unreasonable, as mutually beneficial relations between species can exist, however their study would require a different analytical framework. Yet, if negative values remain rare, it is clear from the self-consistent equation that as tails get thicker, the number of solutions $\mathcal{N}_s$ further decreases before one enters a new regime when negative β 's start proliferating.

3.5 Disorder chaos

Having found that the number of solutions satisfying the non-negativity constraint is near exponential for relevant distributions of β (in the regime $\bar{\beta} > 0$ and $\sigma = \chi/N$), we ask ourselves if we can observe *disorder chaos* in this system. Disorder chaos in this context is essentially the question of the stability of the optimal solution if the disorder β is slightly altered, particularly in the case of large N . Indeed, if there is an exponential number of valid solutions, some with similar values of the objective function, it is not hard to imagine that a slight modification in the disorder could yield a complete reshuffling in the spin configuration. This idea is further supported by the numerical exploration of the solution space that was conducted in Section 3.2.1, where we found that a large number of configurations with little overlap with the optimal solution indeed have very close properties.

This phenomenon has been observed in spin-glasses [142–144], and may be formulated in a formally very similar way. It should be noted that this form of instability under changes in the quenched disorder, sometimes also referred to as static chaos, is not to be confused with temperature chaos [145] (as the use of β might induce some confusion in the spin-glass context). Introducing the perturbation ε , we alter the disorder as

$$\tilde{\beta}_i - 1 = (\beta_i - 1) \left(\frac{1 + \varepsilon v_i}{\sqrt{1 + \varepsilon^2}} \right), \quad (3.52)$$

where v_i is a Gaussian random variable with zero mean and unit variance. This definition allows one to keep the variance of the modified β s unchanged.

To compare the optimal/most diverse non-negative solution to the original problem to the perturbed one, it is necessary to introduce some measure of the overlap between solutions. Recalling our definition of the overlap between two




Figure 3.7: Overlaps obtained numerically for different fixed values of σ , ε while N is varied up to 10^4 , averaged over 64 realizations. Left: $\beta \sim \mathcal{N}(1, \sigma^2)$, data collapsed with $N^{\frac{1}{3}}\sigma^{\frac{1}{3}}\varepsilon$. Right: $\beta \sim \mathcal{U}(1 \pm \sqrt{2}\sigma)$, data collapsed with $N^{\frac{2}{3}}\sigma^{\frac{1}{2}}\varepsilon$. Insets represent the overlap subtracted from 1, its maximum value for identical disorder, plotted in log-log.

configurations, Eq. (3.10), we subsequently define the portfolio correlation as

$$O_N(\mathbf{w}, \tilde{\mathbf{w}}) = \frac{\overline{q_{\{\theta\}, \{\tilde{\theta}\}}(N)}}{\sqrt{\overline{q_{\{\theta\}, \{\theta\}}(N)} \overline{q_{\{\tilde{\theta}\}, \{\tilde{\theta}\}}(N)}}}, \quad (3.53)$$

where $\{\theta\}$ and $\{\tilde{\theta}\}$ correspond to the original and altered configurations respectively. With this definition we ensure $O_N(\mathbf{w}, \tilde{\mathbf{w}}) = 0$ for independent portfolios, and $O_N(\mathbf{w}, \tilde{\mathbf{w}}) = 1$ for $\mathbf{w} = \tilde{\mathbf{w}}$.

The resulting overlap for Gaussian and uniform β 's compatible with market data is shown in Fig. 3.7. Qualitatively, both collapsed plots appear similar, with a decrease in the overlap as N gets large and other parameters are kept fixed. Taking a closer look, it is clear that the disorder chaos is stronger for uniformly distributed β 's, which is consistent with the fact that the number of solutions $\mathcal{N}_s$ is larger in this case, as found in the previous section. While the reduction of the overlap in N is easily understandable, the detailed scaling behavior in σ found numerically is more challenging to describe analytically.

In any case, the logarithmically scaled insets show a clean power law behavior in N . Asymptotically, our results therefore suggest that $O_N(\mathbf{w}, \tilde{\mathbf{w}}) \rightarrow 0$ for any $\varepsilon > 0$ as $N \rightarrow \infty$, characteristic of disorder chaos. On both plots, the data points sampled from empirical β 's are well aligned with numerical simulations, suggesting the chaos observed is robust somewhat beyond the regime studied analytically.

Regardless of the precise behavior of the overlap with the problem's parameters, the disorder chaos observed here is first and foremost a qualitative insight. What this observation tells us has a practical consequence on the way one might approach the systematic construction of a long-only portfolio. Supposing one picks

among 3000 stocks for example, which is a reasonable number for a large asset manager, a change of the order of 10% of the β 's could result in a significant reshuffling in the positions that should be held, particularly if the assets considered are within the same or similar industries for instance. Such a change could e.g. come from empirical estimation errors of the correlations, or simply because the β 's naturally evolve in time depending on the many factors not incorporated in the present risk model. Besides, if one decides to modify the portfolio to match the new optimal result, it is likely that significant transaction costs could come into play, particularly given the highly concentrated nature of large long-only portfolios, so it would rather make sense to choose a portfolio that is a mix of many different quasi-degenerate solutions of the optimisation problem. Note that such a portfolio would then not be optimal in the Markowitz sense (i.e. would not satisfy Eq. (3.2)), but would likely reduce the volatility as the β_i 's are allowed to vary.

From a conceptual point of view, disorder chaos means that two perfectly rational investors with a slightly different method of estimating the β s might end up with very different optimal solutions in the large N limit. As emphasised in [42] and recalled in the introduction of this chapter, the presence of a very large number of quasi-degenerate solutions, at the heart of disorder chaos, leads to some irreducible uncertainty in the decision of agents, even assumed to be fully rational.

In terms of ecological equilibria, this observation also has concrete implications. Indeed, it suggests that a moderate change in the interaction between the N species considered can lead to a significantly different outcome in terms of surviving species at the equilibrium. It seems reasonable to imagine that some physical changes to the environment (e.g. through temperature changes or the introduction of chemicals) could alter the strength of interactions between species, which could then lead to a significantly different equilibrium picture of the ecosystem (on this point, see also [146]).

3.6 Conclusion

Let us summarize what we have achieved in this chapter. Through the introduction of a spin-glass inspired formalism, we have shown that N assets or species can be recombined in an exponential number of solutions satisfying the non-negativity constraint associated to the portfolio and ecological equilibrium problems, in the special case where the interaction matrix is of unit rank. More precisely, we have computed the average (or annealed) number of solutions and have shown that its logarithm grows as N^α , where $\alpha \leq 2/3$ depends on the distribution of asset correlations and interaction strength respectively. This average number does however

not correspond to the typical behavior of the system, observed through a limited number of numerical experiments for example. Indeed, we have found that in this problem the mean number of solutions is heavily skewed by the existence of very unlikely occurrences that yield an exponential number of solutions. In the case of a full rank Gaussian interaction matrix, recent work confirms this idea, as the annealed and quenched complexities of the Lotka-Volterra equilibria are found to be markedly different in some region of the parameter space [147, 148]. Finding the typical (or quenched) number of solutions, by means of a replica calculation, therefore appears to be a natural extension of the present work. We conjecture that the result will be related to the typical sparsity $m(N)$ of the solutions, namely $\log \mathcal{N}_s \propto Nm(N)$. Hence, the number of possible long-only configurations that can be constructed from the N entities considered remains large, specially for a strictly bounded distributions of β 's for which $Nm(N) \sim \sqrt{N}$.

We have also shown numerically that the solution landscape is similar to that of other complex optimisation problems like spin-glasses, i.e. many very different configurations or portfolios are quasi-degenerate, in the sense that they lead to nearly identical values of the objective function (energy for spin-glasses, risk for portfolios). Correspondingly, the phenomenon of “disorder chaos” in spin-glasses, i.e. the extreme sensitivity of the optimal solution on the detailed specification of the problem when N is large, is also present in our long-only portfolio problem (or in its ecological counterpart).

For asset management, this result suggests that, in the presence of transaction costs, the construction of long-only portfolios should account for such an instability and in fact blend together optimal portfolios obtained by slightly varying the risk model (here the value of the β 's). More importantly in the context of this thesis, as emphasized in the original paper from Galluccio *et al.* [42] and recalled above, such a sensitivity is interesting in the sense that it questions the meaning of a rational decision when there is a very large number of quasi-degenerate (or satisficing) solutions. This type of *radical uncertainty* is characteristic of the idea of *radical complexity* at the heart of this thesis.


For ecological equilibria, while there is unfortunately no empirical data to support our choice of interaction matrix and to choose appropriate distributions of β , we believe that most of the conclusions drawn for parameters compatible with stocks should hold for highly competitive environments with a large number of similarly interacting species, as discussed in a different context in [146]. Indeed, the analytical description can be generalized to any values of $\bar{\beta} > 0$ that could be appropriate for the ecology problem, and we have shown that our results are in fact valid for a wide range of distributions of β . While not explicitly discussed here, heterogeneous expected returns (or growth rates) μ_i can be analyzed similarly, see Appendix A.6. We find that the solution is akin to the one obtained with $\mu_i \equiv 1$, with a threshold that is no longer on β alone but on the ratio β/μ .

Chapter 3. Rationality versus complexity: the example of portfolio optimization

In both the portfolio and population dynamics cases, the choice of the effective interaction matrix C_{ij} is the main limiting factor in our study. Extending results to more general (random) matrix models could be an interesting avenue to explore in the future. This being said, the very general formulation of the problem, in essence studying the non-negativity of a linear equation, leads us to believe that long-only portfolios and ecological equilibria are not the only applications for the analytical description detailed in Sec. 3.3. Due to its links with population dynamics, the survival of firms in macroeconomic systems [149,150] could for example be another problem to study with this spin-glass inspired approach.

Key takeaways

- Given N financial assets described by their risk sensitivity to the market, there are on *average* $\sim \exp(N^\alpha)$ number of ways to construct a locally optimal long-only portfolio.
- The *typical* number of configurations is expected to be smaller, although it likely remains very large for plausible values of N .
- A large fraction of these configurations are near optimal in risk despite being composed of entirely different assets.



- As a result of this quasi-degeneracy, a minute change in risk sensitivities leads to a complete reshuffling of the best possible outcome.
- This is an example of *radical complexity* leading to *radical uncertainty*: the very meaning of “rational” decision making is put into question, and we can no longer assume common information.
- The above results can be directly mapped onto the equilibria of Lotka-Volterra equations with self-regulation in population dynamics.

Chapter 3. Rationality versus complexity: the example of portfolio optimization

Chapter 4

Slutsky matrices and the necessity of a global utility

You guys really believe that?

Philip Warren Anderson

In the previous chapter, we have seen how disorder generically gives rise to a large number of nearly equivalent solutions, even in a simple optimization problem, challenging the idea that agents may be entirely rational and have access to identical information from the purely practical point of view. Considering this result, we now turn to a standard economic problem in consumer choice theory, and revisit the classical result after introducing a parameter controlling the effectiveness of the agents' maximization of their utility, i.e. accounting for *bounded* rationality in the decision making process.

This chapter is largely based on [2], written under the supervision of M. Benzaquen and J.-P. Bouchaud, with an enriched discussion on detailed balance violation in the interest of the overall thesis message.

4.1 Consumer choice theory

Consumer choice theory is based on the idea that, for a given bundle of M goods with prices $\mathbf{p} \in \mathbb{R}_+^M$, agents choose a basket $\mathbf{x}$ to maximize their utility function $u(\mathbf{x})$, while subject to the constraint $\mathbf{p} \cdot \mathbf{x} = w$, where w is the consumption budget. Here, the basket compositions x_i , representing the quantity of good i acquired by the agent, are taken to be real numbers, and further assumed to be non-negative.

As touched upon in Sec. 1.2.2, the utility function must then satisfy some elementary conditions. In particular, it is taken to be an increasing function of

Chapter 4. Slutsky matrices and the necessity of a global utility

the quantity of goods for any of the products considered, i.e. $\partial_{x_i} u > 0 \forall i$ and more of anything is always better, with no satiation. While the utility must therefore have a positive first derivative, its rate of increase should be a decreasing function of the absolute quantity of goods. In other words, utility is postulated to be concave, i.e. the marginal utility gain diminishes as goods are accumulated.

Assuming that the agent is fully rational, the optimal basket that maximizes the agent's satisfaction $\mathbf{x}^*$ is then simply given by

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \geq 0} [u(\mathbf{x}) \mid \mathbf{p} \cdot \mathbf{x} = w]. \quad (4.1)$$

Conventionally one defines the *Marshallian* demand $\mathbf{x}(\mathbf{p}, w)$ which corresponds to that obtained from solving problem (4.1), and the *Hicksian* demand (also known as the compensated demand curve) $\mathbf{h}(\mathbf{p}, u)$ defined as the demand that minimizes the expenditure $e(\mathbf{p}, u)$ for a fixed utility level u [151]. Setting $\mathbf{h}(\mathbf{p}, u) = \mathbf{x}^*(\mathbf{p}, e(\mathbf{p}, u))$ and differentiating, one obtains the Slutsky equation describing the change in consumption of good i following a change in the price of good j

$$\frac{\partial x_i}{\partial p_j} = \frac{\partial h_i}{\partial p_j} - x_j \frac{\partial x_i}{\partial w}. \quad (4.2)$$

Changes $\delta \mathbf{x}$ in the optimal basket's composition in response to a price change $\delta \mathbf{p}$ can thus be separated into two contributions: the *substitution effect* (first term in the rhs of Eq. (4.2)), describing how consumption is impacted by changes in relative prices of goods, and the *income effect* (second term in the rhs of Eq. (4.2)), expressing the impact of changes in purchasing power. The substitution effect is often described in terms of the Slutsky matrix $\mathbf{S}$, with entries therefore defined as

$$S_{ij} := \frac{\partial h_i}{\partial p_j} = \frac{\partial x_i}{\partial p_j} + x_j \frac{\partial x_i}{\partial w}. \quad (4.3)$$

Provided that the utility function is sufficiently regular, $\mathbf{S}$ can then be shown to be symmetric, negative semi-definite, and equal to the Hessian of the expenditure function. In practice, the Hicksian demand cannot be observed directly, but the Slutsky matrix can be estimated as the other two terms in Eq. (4.3) should be accessible empirically. In the following, we will give an alternative theoretical expression for the Slutsky matrix in terms of consumption fluctuations.

Before moving on to our statistical physics inspired study of bounded rationality, let us clarify the meaning of the Slutsky matrix in simpler terms, and illustrate the type of situation it describes in practice. Suppose an agent typically buys a quantity x_A of a given consumption good, say apples, and a quantity x_B of another, say bananas. Now, if the price of apples decreases, a kilogramme of apples becomes *relatively cheaper* than that of bananas, so the consumer might be tempted to replace some of its consumption of bananas by apples (after all,

both are tasty fruits). This is the *substitution* effect mentioned above. Simultaneously, the fact that the price of apples has decreased without the overall budget of the agent being altered means that they might have extra change after having satisfied their fruit needs. With this *freed up budget*, the consumer might then be tempted to get some chocolate that they previously could not afford. This is the *income* effect previously introduced. The Slutsky matrix aims at isolating the substitution effect, which is not directly observable as the utility of an agent – that is supposed to be kept fix through this substitution – is not something that one can measure (assuming it even exists).

4.2 Thermodynamics of bounded rationality

As described in the introduction chapter and emphasized above, taking agents to be perfect optimizers seems to be an unrealistic assumption in most contexts, and there are several ways to relax such an assumption. One is that agents have a limited attention and cannot process all the information accessible to them, see e.g. [53] and refs. therein. As a result, agents experience *perceived* prices that differ from real prices, which in turn affects the symmetry of the Slutsky matrix.

Here, we will instead turn to the previously introduced logit stochastic choice rule, such that the stationary choice distribution is given by a Gibbs-Boltzmann measure (see below) with an *intensity of choice* parameter acting as the inverse temperature. As we shall see, this prescription allows us to describe a rather wide range of phenomena while ensuring mathematical tractability of the model (Sec. 2.1.1).

4.2.1 A single agent

Formally, considering first a single agent, one can always write the probability density for selecting the basket of goods $\mathbf{x}$ in the form

$$P(\mathbf{x}) = \begin{cases} \frac{1}{Z} e^{v(\mathbf{x})} & \text{if } \mathbf{p} \cdot \mathbf{x} = w, \\ 0 & \text{otherwise,} \end{cases} \quad (4.4)$$

where $v(\mathbf{x})$ is a certain function overweighting or underweighting basket $\mathbf{x}$, and Z a normalizing factor. Random choice theory postulates that $v(\mathbf{x})$ is related to the utility function through $v(\mathbf{x}) = \beta u(\mathbf{x})$, where we remind that β is known as the intensity of choice. This is the specification we will adopt in the following, however, some of our results below (such as Eq. (4.7)) are in fact more general and would hold for an arbitrary function $v(\mathbf{x})$.

Chapter 4. Slutsky matrices and the necessity of a global utility

We recall that the normalization factor Z is known in statistical physics as the partition function, which, given the hard budget constraint enforced here, writes

$$Z = \int_+ d\mathbf{x} e^{\beta u(\mathbf{x})} \delta(\mathbf{p} \cdot \mathbf{x} - w), \quad (4.5)$$

where $\int_+$ means that we integrate over non-negative baskets $\mathbf{x} \geq 0$. For finite values of β , the basket $\mathbf{x}$ must now be statistically described, since different realizations of the system will lead to different outcomes. A suited definition of the Slutsky matrix must therefore be considered. Here, we propose to replace all x_i by their averages $\langle x_i \rangle$,

$$S_{ij} := \frac{\partial}{\partial p_j} \langle x_i \rangle + \langle x_j \rangle \frac{\partial}{\partial w} \langle x_i \rangle, \quad (4.6)$$

with angular brackets referring to an average over the distribution given in Eq. (4.4), to wit

$$\langle f(\mathbf{x}) \rangle := \frac{1}{Z} \int_+ d\mathbf{x} f(\mathbf{x}) e^{\beta u(\mathbf{x})} \delta(\mathbf{p} \cdot \mathbf{x} - w).$$

Note that in the limit $\beta \rightarrow \infty$, one recovers the standard Slutsky matrix since $\langle x_i \rangle \rightarrow x_i^*$.

Our set up of the problem allows us to draw several analogies with the statistical mechanics foundation of thermodynamics introduced in Chap. 2. For example, the strict application of the budget constraint $\delta(\mathbf{p} \cdot \mathbf{x} - w)$ is reminiscent of the so-called *canonical* ensemble, where the conservation of the number of particles in the system is strictly enforced. We will see later that one can also work in the analogue of the *grand-canonical* ensemble where the budget constraint is only enforced on average. This eases some analytical calculations while being equivalent to the canonical ensemble in some limits (for example when the number of goods is large). One can argue that in some cases, allowing the budget to fluctuate (due to loans for example) can be realistic as well.

4.2.2 A fluctuation-response relation

More interestingly, statistical mechanics also provides relations between the response of certain quantities to external perturbations to spontaneous fluctuations of these quantities in the absence of perturbations. These relations can be established using the derivatives of the partition function, assuming an equilibrium state has indeed been reached. In particular, the Slutsky matrix can indeed be expressed in terms of other correlations (as was first mentioned in [152]). In App. B.1.1, we derive the following “fluctuation-response” formula in the single agent case:

$$S_{ij} = -\Gamma \langle x_i x_j \rangle_c - \partial_w \langle x_i x_j \rangle_c, \quad (4.7)$$

with $\Gamma = \partial_w \log Z$ and the “connected” correlations $\langle x_i x_j \rangle_c := \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$. Note that, interestingly, the utility function $u(\mathbf{x})$ does not appear explicitly in Equation (4.7), which is manifestly symmetric in i, j . It shows that even with bounded rationality, the Slutsky matrix is still symmetric, for any value of β , not only in the rational limit $\beta \rightarrow \infty$. Hence, *the symmetry of the Slutsky matrix may not be used as an argument for or against the rationality of economic agents*, contrarily to some claims made in [152] for example.

Our fluctuation formula Eq. (4.7) is also interesting from an econometric standpoint, as it provides a way to measure the Slutsky matrix without varying prices. Measuring response quantities from equilibrium correlations is in fact commonly used in statistical mechanics, through what is referred to as *fluctuation-dissipation* relations [153], or, in a restricted context, to Einstein’s relation relating mobility to diffusion for Brownian particles. Note that we do not have to specify the nature of the dynamics or the structure of temporal shocks, provided these lead to an equilibrium of the form given by Eq. (4.4).

One can go one step further and eliminate all derivatives from Eq. (4.7), but only provided the utility function is known. One finds (see App. B.1.3)

$$\begin{aligned} \frac{\partial}{\partial w} \langle x_i x_j \rangle_c = \frac{\beta}{p_k} & \left[\left\langle x_i x_j \frac{\partial u}{\partial x_k} \right\rangle + 2 \langle x_i \rangle \langle x_j \rangle \left\langle \frac{\partial u}{\partial x_k} \right\rangle - \langle x_i x_j \rangle \left\langle \frac{\partial u}{\partial x_k} \right\rangle \right. \\ & \left. - \langle x_i \rangle \left\langle x_j \frac{\partial u}{\partial x_k} \right\rangle - \langle x_j \rangle \left\langle x_i \frac{\partial u}{\partial x_k} \right\rangle \right], \end{aligned} \quad (4.8)$$

and

$$\Gamma = \frac{\beta}{p_k} \left\langle \frac{\partial u}{\partial x_k} \right\rangle, \quad (4.9)$$

both equations being valid for an arbitrary choice of k . These expressions considerably simplify in the near-rational limit $\beta \rightarrow \infty$, see Section 4.2.5 below.

4.2.3 Many agents

In practice, agents make correlated choices and we must adapt our formalism to treat interactions. We thus consider N agents, indexed by $\alpha = 1, \dots, N$, and postulate that the stationary distribution for a simultaneous set of choices $\{\mathbf{x}^\alpha\}$ is still given by a constrained Gibbs-Boltzmann distribution

$$P(\{\mathbf{x}^\alpha\}) = \begin{cases} \frac{1}{Z} e^{\beta U(\{\mathbf{x}^\alpha\})} & \text{if } \mathbf{p} \cdot \mathbf{x}^\alpha = w^\alpha, \forall \alpha, \\ 0 & \text{otherwise,} \end{cases} \quad (4.10)$$

where $U(\{\mathbf{x}^\alpha\})$ is a certain function that generalizes the single agent utility function while encoding interactions between agents. Interactions mean that U cannot in general be written as a sum of individual utility functions u^α . At this stage, it is

Chapter 4. Slutsky matrices and the necessity of a global utility

important to note that while this function should always exist in the steady-state, explicitly relating it to dynamics at the agent level may not be straightforward, see Section 4.5.1. Note in particular that Eq. (4.10) does *not* require the existence of a “central planner”.

The probability that the N agents choose a bundle of goods $\{\mathbf{x}^\alpha\}$ is proportional to $\exp(\beta U)/Z_N$, where the aggregate partition function Z_N writes

$$Z_N = \int_+ \left(\prod_{\alpha=1}^N \prod_{i=1}^M dx_i^\alpha \right) e^{\beta U(\{\mathbf{x}^\alpha\})} \prod_{\alpha=1}^N \delta(\mathbf{p} \cdot \mathbf{x}^\alpha - w^\alpha). \quad (4.11)$$

Here, we integrate over the $M \times N$ degrees of freedom, i.e. the quantities x_i^α of good i consumed by agent α , while enforcing that all agents respect their own specific budget w^α . Analytically computing this partition function is usually very difficult due to the product of Dirac δ distributions in the integrand. Depending on the form of the utility, one might need to slightly relax the budget constraints. As mentioned above, a way to do this is to move to the so-called *grand-canonical* partition function $\mathcal{Z}_N$, defined as

$$\mathcal{Z}_N = \int_+ \left(\prod_{\alpha=1}^N \prod_{i=1}^M dx_i^\alpha \right) e^{\beta [U(\{\mathbf{x}^\alpha\}) - \sum_{\alpha} \mu^\alpha \mathbf{p} \cdot \mathbf{x}^\alpha]}, \quad (4.12)$$

where the μ^α , known in physics as “chemical potentials”, are fixed by enforcing that budgets are satisfied on average, i.e. $\mathbf{p} \cdot \langle \mathbf{x}^\alpha \rangle = w^\alpha$, $\forall \alpha$. In general, the two partition functions are not equivalent; however, many quantities calculated from the two ensembles become identical in the large M limit and/or in the rational limit $\beta \rightarrow \infty$, see Section 4.5.2 for a detailed discussion.

Consistent with the single agent definition, we take the individual Slutsky matrix of agent α to be given by

$$S_{ij}^\alpha := \frac{\partial}{\partial p_j} \langle x_i^\alpha \rangle + \langle x_j^\alpha \rangle \frac{\partial}{\partial w^\alpha} \langle x_i^\alpha \rangle, \quad (4.13)$$

where we have assumed that prices are the same for all agents. Note that we have also taken a uniform system-wide rationality parameter β , although a generalization to different β^α ’s is possible and would be an interesting extension of our work.

As in the single agent case, the partition function allows one to derive a fluctuation-response expression for the Slutsky matrix in terms of correlations. Equation (4.13) may be rewritten (see App. B.1.2) in what will be referred to as its “thermodynamic” form

$$S_{ij}^\alpha = - \sum_{\gamma} \left[\Gamma_{\gamma} \langle x_i^\alpha x_j^\gamma \rangle_c + \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle_c + (1 - \delta_{\alpha\gamma}) \langle x_j^\gamma \rangle \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha \rangle \right], \quad (4.14)$$

with $\Gamma_\gamma = \partial_{w^\gamma} \log Z_N$. This is the central theoretical result of this chapter, which, to the best of our knowledge, is new. Again, the function $U(\{\mathbf{x}^\alpha\})$ does not explicitly appear in this result. Importantly, unlike in the single-agent case, Eq. (4.14) does not allow one to infer anything general about the symmetry of the Slutsky matrices. In the case where interactions between agents are negligible, i.e. when $U(\{\mathbf{x}^\alpha\})$ is the sum of individual utility functions, correlations between agents are zero whenever $\gamma \neq \alpha$ and we recover the single agent expression, as expected.

Assuming the function $U(\{\mathbf{x}^\alpha\})$ is known, one can again go further and express derivatives with respect to budgets w^γ as a function of some correlations. For our problem we find a generalization of the formula obtained for a single agent in the canonical ensemble (see App. B.1.3)

$$\begin{aligned} \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle_c = \frac{\beta}{p_k} & \left[\left\langle x_i^\alpha x_j^\gamma \frac{\partial U}{\partial x_k^\gamma} \right\rangle + 2 \langle x_i^\alpha \rangle \langle x_j^\gamma \rangle \left\langle \frac{\partial U}{\partial x_k^\gamma} \right\rangle - \langle x_i^\alpha x_j^\gamma \rangle \left\langle \frac{\partial U}{\partial x_k^\gamma} \right\rangle \right. \\ & \left. - \langle x_i^\alpha \rangle \left\langle x_j^\gamma \frac{\partial U}{\partial x_k^\gamma} \right\rangle - \langle x_j^\gamma \rangle \left\langle x_i^\alpha \frac{\partial U}{\partial x_k^\gamma} \right\rangle \right], \end{aligned} \quad (4.15)$$

as well as

$$\Gamma_\gamma = \frac{\beta}{p_k} \left\langle \frac{\partial U}{\partial x_k^\gamma} \right\rangle, \quad (4.16)$$

both again valid for any good k . From these expressions and given a utility function $U(\{\mathbf{x}^\alpha\})$, all terms in Eq. (4.14) can be computed in principle, at least numerically. Bear in mind, however, that these relations are only valid if the system has reached equilibrium, which might take a very long time, for example near a phase transition point. Note here the important distinction between the *thermodynamic* equilibrium that we are referring to, which is characterized by the fact that the system correctly samples the Gibbs-Boltzmann distribution of Eq. (4.10), and the idea of an *economic* equilibrium that is postulated from the very beginning. As we shall see in Section 4.5.1, one can easily construct a system that still represents an economic equilibrium but is devoid of a thermodynamic equilibrium state.

4.2.4 Aggregate Slutsky matrices

There are *a priori* two possible definitions of the Slutsky matrix at the aggregate level. One is simply to take the average over all agents of the individual Slutsky matrices, to wit

$$\bar{S}_{ij} := \frac{1}{N} \sum_{\alpha} S_{ij}^{\alpha}. \quad (4.17)$$

Chapter 4. Slutsky matrices and the necessity of a global utility

However, empirical measurements often rely on estimates at the aggregate level. In that case, a better suited definition uses aggregate consumption:

$$\mathcal{S}_{ij} := \frac{\partial}{\partial p_j} \langle \bar{x}_i \rangle + \langle \bar{x}_j \rangle \frac{\partial}{\partial \bar{w}} \langle \bar{x}_i \rangle, \quad (4.18)$$

with overlines indicating averaging over agents e.g.

$$\bar{x}_i := \frac{1}{N} \sum_{\alpha} x_i^{\alpha},$$

and similarly for the average budget $\bar{w}$. In that case, as shown in App. B.1.4, the thermodynamic expression becomes

$$\mathcal{S}_{ij} = -\frac{1}{N} \sum_{\alpha, \gamma} \left[\Gamma_{\gamma} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c + \frac{\partial}{\partial w^{\gamma}} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c + \left(\langle x_j^{\gamma} \rangle - \kappa^{\gamma} \bar{x}_j \right) \frac{\partial}{\partial w^{\gamma}} \langle x_i^{\alpha} \rangle \right], \quad (4.19)$$

with the possibly heterogeneous factor

$$\kappa^{\gamma} := \frac{\partial w^{\gamma}}{\partial \bar{w}}.$$

Clearly, if consumption is proportional to wages and if all wages scale with the average wage, i.e. $w^{\gamma} = \kappa^{\gamma} \bar{w}$, then $\langle x_j^{\gamma} \rangle = \kappa^{\gamma} \bar{x}_j$, and there is no contribution from the last term. In this case, we find an expression very close to the single-agent case. More generally, $\mathcal{S}_{ij}$ has no reason to be symmetric, except when all agents are identical. In such a case, even in the presence of interactions and for bounded rationality ($\beta < +\infty$), $\mathcal{S}_{ij}$ is always symmetric, whereas $\bar{\mathcal{S}}_{ij}$ is not, as we in the next sections.

4.2.5 Near-rational limit $\beta \rightarrow \infty$

In order to simplify the problem and get some intuition, we place ourselves in the near-rational case where $\beta \rightarrow \infty$. This corresponds to the low temperature case in a physical system, where one can expect that all relevant configurations $\{\mathbf{x}^{\alpha}\}$ are close to the optimal configuration $\{\mathbf{x}^{\alpha*}\}$ that maximizes the function $U(\{\mathbf{x}^{\alpha}\})$ subject to budget constraints. We thus write $\delta x_i^{\alpha} := x_i^{\alpha} - x_i^{\alpha*}$ and Taylor-expand to second order, resulting in

$$U(\{\mathbf{x}^{\alpha}\}) \approx U^* + \frac{1}{2} \{\delta \mathbf{x}^{\alpha}\}^{\top} \mathbf{H} \{\delta \mathbf{x}^{\alpha}\} + O(\delta x^3), \quad (4.20)$$

with $\mathbf{H}$ the $(M \times N) \times (M \times N)$ Hessian of the system evaluated at the maximum of the function $U(\{\mathbf{x}^{\alpha}\})$, for a given set of budget constraints w^{α} . Here we only consider deviations $\{\delta \mathbf{x}^{\alpha}\}$ that all satisfy the budget constraints, so that the first

derivative terms U' are zero as we expand around a maximum along all directions but one. Calculating the partition function and correlations then simply amounts to computing Gaussian integrals, although the budget constraints must still be handled with care. At this stage we emphasize that in the context of large β all fluctuations are small and Gaussian, and therefore our results in that limit are actually very general. By taking the Fourier representation of the Dirac δ (see App. B.1.5), we finally find, to leading order in β^{-1} :

$$C_{ij}^{\alpha\gamma} := \langle x_i^\alpha x_j^\gamma \rangle_c = -\frac{1}{\beta} \left[(\mathbf{H}^{-1})_{ij}^{\alpha\gamma} - \sum_{\eta, \nu, k, \ell} (\mathbf{H}^{-1})_{ik}^{\alpha\eta} p_k (\mathbf{G}^{-1})^{\eta\nu} p_\ell (\mathbf{H}^{-1})_{\ell j}^{\nu\gamma} \right], \quad (4.21)$$

with the $N \times N$ matrix $\mathbf{G}$ defined as

$$G^{\eta\nu} = \sum_{i,j} p_i (\mathbf{H}^{-1})_{ij}^{\eta\nu} p_j, \quad (4.22)$$

and where the second term in the rhs of Eq. (4.21) is the result of the constraint being applied.

Let us first illustrate this formula in the $N = 1$ case. In the canonical ensemble and large β regime, one has

$$\left\langle \frac{\partial u}{\partial x_k} \right\rangle = \lambda p_k + O(\beta^{-2}),$$

where λ is the Lagrange parameter enforcing that the budget constraint is satisfied. Hence the first term Eq. (4.7) remains finite since Γ diverges as β when $\langle x_i x_j \rangle_c$ tends to zero as β^{-1} . The second term, on the other hand, tends to zero at least as β^{-1} . It then follows that for a single and near-rational agent,

$$\mathbf{S} = \lambda \left[\mathbf{H}^{-1} - \mathbf{u} \mathbf{u}^\top \right] + O(\beta^{-1}), \quad \mathbf{u} = \frac{\mathbf{H}^{-1} \mathbf{p}}{\sqrt{\mathbf{p}^\top \mathbf{H}^{-1} \mathbf{p}}}. \quad (4.23)$$

This is the classic expression for the Slutsky matrix, which our framework therefore allows to recover in the corresponding limit, with corrections in β^{-1} that can be computed. Since the Hessian is both symmetric and negative semi-definite at a utility maximum, we recover the classic properties of the Slutsky matrix. “Homogeneity” is also easily recovered by checking that multiplying the matrix by $\mathbf{p}$ indeed gives a zero eigenvalue.

When $N > 1$, we will need to specify the function $U(\{\mathbf{x}^\alpha\})$ to make the final result more explicit. Keeping U fully general and taking the limit $\beta \rightarrow \infty$ only allows one to simplify the general expression Eq. (4.14) to

$$S_{ij}^\alpha = -\beta \sum_{\gamma} \lambda_\gamma C_{ij}^{\alpha\gamma} - \sum_{\gamma \neq \alpha} \frac{\langle x_j^\gamma \rangle}{p_j} \sum_{k, \nu} H_{ik}^{\alpha\nu} C_{ik}^{\alpha\nu} + O(\beta^{-1}), \quad (4.24)$$

where we have used Eq. (B.20) and expanded $\partial_x U$ to first order in δx . A similar expression can be derived for the aggregate Slutsky matrix $\mathcal{S}_{ij}$ as well. Remember that $\mathbf{C}$ is of order β^{-1} (see Eq. (4.21)) so S_{ij}^α is well behaved in the large β limit. Although the final expression is not transparent, it is clear that S_{ij}^α has no reason to be always symmetric, except when agents' choices are uncorrelated, in which case $C_{ij}^{\alpha\gamma} = 0$ whenever $\alpha \neq \gamma$ and $C_{ij}^{\alpha\alpha}$ is symmetric by construction. We will now turn to an explicit model with herding, where the asymmetry of Slutsky matrix can be made apparent.

4.3 Animal spirits

Both anecdotal evidence about fads and fashion and more serious scientific studies point to the fact that agents' choices can be strongly influenced by the choice of others (see e.g. [15, 154–158]), an effect also known as “keeping up with the Joneses” [159, 160], see also [161]. We now propose a family of models which account for interactions between boundedly rational agents, allowing for strong imitation and herding, or “animal spirits” as famously coined by J.M. Keynes [162]. Indeed, these models lead, in some regions of parameters, to “concentration” (or “condensation”) of choices, much in the spirit of the model proposed by Borghesi and Bouchaud [163] (see also [164] for a recent extension). As we shall see, close to the concentration transition, the non-symmetric contribution to the Slutsky matrix reaches a maximum.

4.3.1 Interactions and herding

In order to study imitation or fashion effects, we first consider agents with log-utilities and take the aggregate function $U(\{\mathbf{x}^\alpha\})$ to simply be the sum of all agent-specific utilities,

$$U(\{\mathbf{x}^\alpha\}) = \sum_{\alpha=1}^N \sum_{i=1}^M a_i^\alpha \log x_i^\alpha. \quad (4.25)$$

where a_i^α describes the preference of agent α for good i . In the following, we assume that agents are homogeneous ($a_i^\alpha = a_i$, $\forall \alpha$) but interacting, by which we mean that the preference for good i depends on how much good i is consumed by other agents. Mathematically, we posit that the preference for goods increases with the k -th power of their average consumption:

$$a_i \rightarrow a_i \left[1 + c(\bar{x}_i)^k \right], \quad (4.26)$$

where we remind $\bar{x}_i$ is the average consumption of good i (over all agents), and c and k are non-negative parameters that describe the strength and nature of the

interactions.²³ The interpretation of this specific aggregate utility-like quantity at the agent level is discussed in Section 4.5.1.

The non-interacting case $c = 0$ (or equivalently $k = 0$, up to some rescaling of the a_i) may be treated exactly in the canonical ensemble, i.e. strictly enforcing the budget constraint. As detailed in App. B.3.1, the equilibrium configurations are given by

$$\langle x_i^\alpha \rangle = \frac{w^\alpha}{p_i} \frac{1 + \beta a_i}{\sum_k (1 + \beta a_k)}, \quad (4.27)$$

which will henceforth be referred to as the non-condensed or uniform solution. This solution matches results from numerical experiments, as illustrated in the two top panels of Figure 4.1(a) in the fully rational case, for example. For details about how numerical simulations have been performed, see App. B.2. The agent specific Slutsky matrices S_{ij}^α can then also be written explicitly from the original definition, verifying both symmetry and negative semi-definiteness for any β . In this non-interacting case, budget heterogeneities simply affect the magnitude of any given agent Slutsky matrix entries and are thus inconsequential for the properties of interest.

Taking $k > 0$ and increasing c , we expect the system to progressively depart from this solution and concentrate on some product(s) as $c \rightarrow \infty$, as observed in the numerical simulations shown in Fig. 4.1(a), bottom panel. To evaluate how this concentration occurs, we start by taking $w^\alpha = w$, $\forall \alpha$, that is a system of interacting but identical agents, both in wealth and preferences. As is often the case in statistical mechanics, the partition function cannot be computed exactly for interacting systems for general N , but can be more and more accurately approximated in the large N limit.

In our case, it is convenient to relax the budget constraint and to place ourselves in the grand-canonical ensemble where the budget constraint is only enforced on average. The procedure, detailed in App. B.3.2, allows us to rewrite the grand-canonical partition function as an integral over the mean consumption vector $\bar{\mathbf{x}}$

$$\mathcal{Z}_N = \int_0^\infty d\bar{\mathbf{x}} e^{-N\beta f(\bar{\mathbf{x}})}, \quad (4.28)$$

where f is usually called the “free energy density”, here given by

$$\begin{aligned} \beta f(\bar{\mathbf{x}}) = & \sum_i [\beta \mu p_i \bar{x}_i - (1 + \beta a_i [1 + c(\bar{x}_i)^k]) (1 + \log \bar{x}_i - \log(1 + \beta a_i [1 + c(\bar{x}_i)^k])) \\ & - \log \Gamma(1 + \beta a_i [1 + c(\bar{x}_i)^k])] + o(1). \end{aligned} \quad (4.29)$$

with μ the “chemical potential” that can be thought of as the Lagrange multiplier used to enforce the averaged constraint $\sum_i \bar{x}_i p_i = w$, taken to be identical for

²³For an alternative specification, see App. B.4.

all agents, and $\Gamma(\cdot)$ is the Gamma function. Importantly, this grand-canonical description allows one to have a free energy density $f(\bar{\mathbf{x}})$ that is a sum over entirely decoupled goods. The only coupling is through the value of μ , which can generally not be expressed as an explicit function of $\bar{x}_i$ and must therefore be determined self-consistently.

Given that $N \rightarrow \infty$, the shape of the free energy then completely determines the state of the system. Indeed, for large N the partition function can be estimated using Laplace's method, such that the values $\bar{x}_i^* = \langle x_i \rangle$ that minimize the free energy are overwhelmingly more probable than any other values. Setting $c = 0$ and solving the set of equations $\frac{\partial f}{\partial \bar{x}_i} = 0$, one can for example check that the previously obtained solution is recovered.

As in statistical physics, we expect to identify the phase transition from the uniform to the *condensed* (concentrated) phase where herding dominates [15, 163]. This occurs when the single minimum in free energy associated to the non-condensed solution becomes a maximum while one or several new minima appear. Such a change of topology occurs for some value c_{crit} , as illustrated in Fig. 4.1(b) in the case where there are only two products. For given values of c , the depth of the minima in the concentrated phase will depend on the a_i and p_i . The most favorable configuration (in this case having more of the least expensive of the two products, as here $a_i = 1, \forall i$) is associated to a lower free energy.

4.3.2 Concentration for $\beta \rightarrow \infty$

In order to find precisely when and how concentration occurs, we first study the case of fully rational agents, $\beta \rightarrow \infty$. Carefully rescaling the free energy density, we find that an extremum is reached for the configurations $\bar{\mathbf{x}}^*$ satisfying

$$a_i[1 + c(\bar{x}_i^*)^k(1 + k \log \bar{x}_i^*)] = \mu p_i \bar{x}_i^*, \quad (4.30)$$

with the value of the chemical potential being such that

$$\sum_i \bar{x}_i^* p_i = w. \quad (4.31)$$

As previously described, the critical value $c_{\text{crit}} = c_\infty$ where this transition occurs in the fully rational case can be found by looking at the Hessian of the free energy evaluated at the non-condensed solution. Doing so (see App. B.3.2), and choosing for the sake of simplicity $a_i = p_i = 1$ for all agents, one finds

$$\frac{1}{c_\infty} = \left(\frac{w}{M}\right)^k \left[2k - 1 + k(k - 1) \log \left(\frac{w}{M}\right)\right]. \quad (4.32)$$

As long as the right hand side is strictly positive, there will therefore be a value of c above which concentration occurs when $\beta \rightarrow \infty$, while if this is not the




Figure 4.1: (a) Monte Carlo simulations at $\beta \rightarrow \infty$ for $M = 6$ products, $N = 256$ identical agents, $k = 2$ and $p_i = a_i = 1$, compared to the non-condensed solution (dashed line). Top: Non-Condensed (NC) region where the value of c_∞ diverges [$w = 0.5$, $c = 1$]. Middle: Possible Concentration (PC) region for $c < c_\infty$ [$w = 10$, $c = 0.01$]. Bottom: PC region for $c > c_\infty$ [$w = 10$, $c = 0.1$], the system concentrates on one of the products and departs from the non-condensed solution. (b) Illustration of the free energy f for some rationality β in the PC phase for $M = 2$ products, $p_2 > p_1$ and $a_2 = a_1$. Left: before the transition, $c < c_{\text{crit}}$. Right: after the transition $c > c_{\text{crit}}$, the deeper minima corresponding to increasing values of c . (c) Theoretical phase diagram of the NC and PC regions at $\beta \rightarrow \infty$, $a_i = p_i = 1$.

case concentration never occurs, regardless of the strength of interactions. The resulting theoretical phase diagram is shown in Fig. 4.1(c) for $k = 2$, a value that we shall keep fixed henceforth. This phase diagram is in perfect agreement with our numerical simulations, as shown in Fig. 4.1(a). Note that this procedure can be repeated for different values of a_i and p_i , leading to qualitatively similar results. Since the Hessian of the free energy is diagonal in our case, the critical value c_∞ would then correspond to the first change of sign of a diagonal element of the matrix.

4.3.3 Finite β effects

Placing ourselves in the regions where condensation does occur in the fully rational limit, we now set out to understand how bounded rationality might alter the phase transition. In the general case, analytical expressions are difficult to obtain. However, numerically finding where the Hessian (which is still diagonal in i, j) loses stability for the non-condensed solution yields a semi-analytical critical line in (c, β) space. It should be noted that this condition, explicitly given in App. B.3.2, is independent of the value of μ , suggesting that the location of the transition is identical in the canonical and grand-canonical descriptions.

Our analytical result can be compared to numerical simulations, for which the transition to fashion dominated consumption is identified by looking at the rescaled Herfindahl index

$$\tilde{H} = \frac{\sum_i (\bar{x}_i p_i / w)^2 - 1/M}{1 - 1/M}.$$

This index takes the values 0 and 1 in the uniform ($\bar{x}_i = w/(Mp_i)$, $\forall i$) and fully concentrated ($\bar{x}_i = w/p_i$ for one product and $= 0$ for all others) cases respectively. As shown in Fig. 4.2(a), the numerical phase diagram and the theoretical critical line match very well, despite the fact that the semi-analytical calculation is based on the grand-canonical ensemble, whereas numerical simulations strictly enforce the budget constraints for all agents. Furthermore, using the theoretical values for $c_{\text{crit}}(\beta)$ to rescale the evolution of the average basket as a function of c/c_{crit} (as plotted in Fig. 4.2(c)), we observe that the evolution of the mean basket and related quantities appears to be largely independent of β .

4.4 Consequences on the Slutsky matrix

Using a simple interacting model, we have shown that introducing herding in the problem leads to a concentration transition and radical changes in the way agents allocate their budget among the M available goods. We now set out to




Figure 4.2: Effect of bounded rationality on concentration from numerical experiments for $M = 4$ products, $N = 64$ agents, $k = 2$, $w = 10$, $p_i = a_i = 1$. (a) Herding phase diagram, with the normalized Herfindahl index $\bar{H}$ calculated over the mean basket among agents, giving $\bar{H} = 0$ in the Non-Condensed (NC) phase and quickly reaching $\bar{H} = 1$ when condensation (C) occurs, indicating fashion dominated consumption. The critical line calculated analytically is shown by the dashed line, and perfectly matches numerical results. (b) Monte-Carlo dynamics of the mean basket for $\beta = 1$, $c \approx c_{\text{crit}}$, showing large fluctuations and switching behavior. The non-condensed solution is shown by the dashed line. (c) Evolution of the rescaled average basket for different values of the rationality parameter β as a function of c , collapsed using the analytical values of c_{crit} visible in the phase diagram. The analytical solution at $\beta \rightarrow \infty$ is shown as the continuous line. Inset: normalized Herfindahl, that clearly reveals the transition around c_{crit} for all values of β .

evaluate the impact of such a transition on the Slutsky matrix and its properties, in particular its negative semi-definiteness and symmetry.

Whereas the grand-canonical theory allowed us to calculate $\langle x_i^\alpha \rangle = \bar{x}_i^*$ for the entire range of β and c self-consistently, the absence of an explicit expressions prevents us from directly computing S_{ij}^α . Instead, because we have found that results are largely independent of β provided c is rescaled as c/c_{crit} , we may gain insight from the Gaussian approximation of the Slutsky matrix introduced in Section 4.2.3 and valid for $\beta \rightarrow \infty$. Equation (4.24) can now be made explicit and writes (see App. B.3.3)

$$S_{ij}^\alpha = -\beta \sum_{\gamma} \left(\frac{a_\ell}{p_\ell \bar{x}_\ell^*} [1 + c(\bar{x}_\ell^*)^k (1 + k \log \bar{x}_\ell^*)] + (1 - \delta_{\alpha\gamma}) \frac{a_j}{p_j \bar{x}_j^*} [1 + c(\bar{x}_j^*)^k] \right) C_{ij}^{\alpha\gamma}, \quad (4.33)$$

for any $\ell = 1 \dots M$. As the most probable values $\bar{\mathbf{x}}^*$ can be obtained by solving Eq. (4.30), the last step is to invert the Hessian of the function $U(\{\mathbf{x}^\alpha\})$ in order to compute the covariance matrix $\mathbf{C}$. Due to the homogeneous nature of the interactions, the Hessian has a very regular structure and may thus be inverted explicitly, as detailed in App. B.3.3. Doing so and replacing in Eq. (4.21), we find that $\mathbf{C}$ has the following form

$$C_{ij}^{\alpha\gamma} = \frac{1}{\beta} \left[\varphi_{ij} \delta_{\alpha\gamma} + \frac{1}{N} \psi_{ij} \right],$$

Chapter 4. Slutsky matrices and the necessity of a global utility

where φ_{ij} and ψ_{ij} are both symmetric and $O(1)$ in N , and, of course, depend on $\bar{\mathbf{x}}^*$. Such correlations, which can be checked to match numerical simulations very well, finally yield an identical Slutsky matrix for all agents,

$$S_{ij} = -\frac{a_j}{p_j \bar{x}_j^*} [kc(\varphi_{ij} + \psi_{ij})(\bar{x}_j^*)^k \log \bar{x}_j^* + [1 + c(\bar{x}_j^*)^k] \varphi_{ij}]. \quad (4.34)$$

Together with Eq. (4.14), this is a key result of our work.

Thus far, we have considered the case where all agents and goods are identical. In the present context, however, it is more illustrative to break the symmetry between products by introducing heterogeneous p_i and a_i , leading to clear preferences between products. An example with $M = 4$ goods is given in Fig. 4.3, showing a very good agreement between the fully rational theory and numerical simulations with finite β and N sufficiently far from the transition region.²⁴ Upon inspection of opposing entries, it quickly appears that the symmetry of the matrix is not satisfied near the critical value of c where condensation first occurs. Note that all matrix entries quickly become very small once the system has entered the concentrated phase. This can be understood for large c by making the *ansatz* $\bar{x}_i^* = w/p_i - \frac{1}{c} \sum_{j \neq i} y_j$ for the dominant product and $\bar{x}_j^* = y_j/c$ for others. Plugging such a guess in Eq. (4.30) indeed solves the equations, and predicts $S_{ij} \sim 1/c$ to leading order in c^{-1} .

The first property of the matrix that interests us is its spectrum, and in particular the non-positivity of its eigenvalues. As shown in Fig 4.4(a), the fully rational theory provides a very good description of the matrix eigenvalues, which remain non-positive for the entire range of c . The leading eigenvalue actually peaks close to the transition. Consistent with the decay of the matrix entries themselves, the magnitude of the eigenvalues also vanish as c increases beyond the transition. As a result, our theory and numerical experiments show that the Slutsky matrix does remain negative semi-definite for all value of c and β . The main consequence of the herding transition on the spectrum is the decay in the magnitude of eigenvalues, as confirmed by looking at the trace of the matrix shown in Fig. 4.4(b).

The other essential property of the Slutsky matrix is its symmetry. Due to the decay in the magnitude of the matrix entries once the system has entered the herding phase, this property becomes difficult to measure as c is increased beyond c_{crit} . Indeed, as both S_{ij} and S_{ji} become very small, any finite numerical error ε affecting either entries will result in a very large relative asymmetry, at which point most common measures of asymmetry will fail. To minimize the impact of

²⁴Close to the transition, agents flip-flop between different basket compositions (Fig. 4.2(b)), which leads to anomalously strong fluctuations and corrections to the Laplace saddle point method used in all our analytical calculations. This is a well known effect in statistical physics, which leads to interesting phenomena in their own right, but that we do not explore further here.




Figure 4.3: Evolution of the individual Slutsky matrix entries with c for $M = 4$ product, $k = 2$, $w = 10$, $\mathbf{p} = [2.2, 2.1, 1.6, 2.3]$, $a_i = 1$ from theory at $\beta \rightarrow \infty$ (continuous lines) and numerical experiments for $\beta = 4$, $N = 16$, calculated using both the pathwise derivative estimates introduced App. B.2 (circles) and the fluctuation-response relations (crosses), averaged over all agents, errorbars indicating one standard deviation. The dashed horizontal line indicates 0, while dotted lines correspond to ± 0.1 as no vertical scale is shown for visual clarity. Opposing entries have identical vertical scales as to highlight the strong asymmetry of some entries (e.g. 13 and 31) close to the herding transition.




Figure 4.4: Properties of the individual Slutsky matrix shown in Fig. 4.3. (a) Real part of the Slutsky matrix eigenvalues λ_i . (b) Trace of the matrix, dotted line indicating zero. (c) Asymmetry measure χ . The analytical results for the aggregate matrix $\mathcal{S}$ are given by the dashed lines for comparison (note that $\mathcal{S}$ remains symmetric for all c).

such numerical errors on the conclusions of our study, we propose an asymmetry measure χ defined as

$$\chi = \left| \frac{\sum_{\alpha} \sum_{j < i} (S_{ij}^{\alpha} - S_{ji}^{\alpha})}{\sum_{\alpha} \sum_{j < i} (S_{ij}^{\alpha} + S_{ji}^{\alpha})} \right|, \quad (4.35)$$

which should be equal to zero for symmetric matrices, and diverge in the anti-symmetric case. Employing this metric with the $\beta \rightarrow \infty$ theoretical Slutsky matrix, we find that the matrix is very close to (but not exactly) symmetric far from the condensation transition. In the vicinity of c_{crit} however, strong interactions give rise to a significant value of χ , contradicting the conventional lore even in the fully rational case. This theoretical result, compared with numerical results, is shown in Fig. 4.4(c), while also visible in Fig. 4.3. Note that, as expected given our choice of identical agents with identical budgets, the previously introduced *aggregate* Slutsky matrix $\mathcal{S}$ remains symmetric even as the concentration of choice occurs. The theoretical results for this alternate definition are shown by the dashed lines in Fig. 4.4, where it is clear that the asymmetry measure χ is always zero. Interestingly, the eigenvalues of the matrix are very similar for the individual (S_{ij}) and aggregate ($\mathcal{S}_{ij}$) definitions.

Regardless of the metric, Fig. 4.4 also illustrates the discrepancy between the equilibrium theory we have devised and the numerical measurements in the transition region. As expected, the system indeed takes a very long time to reach the Gibbs-Boltzmann distribution when the transition occurs. These non-equilibrium effects are also reflected in the difference between the numerical measurements obtained with finite differences and those calculated using the fluctuation-response relations and the associated “thermodynamic” expression of the Slutsky matrices. We expect that such effects will also be present in real empirical data, specially if herding effects bring the system close to a transition point, as seemed to be the

case in the Salganik *et al.* experiment [157, 163].

4.5 Discussion

Before concluding our study, let us examine two subtle points that the above analysis has treated in a somewhat cavalier way.

4.5.1 Global vs. individual utilities

In the numerical (Monte Carlo) investigation of our model of interacting agents, we have assumed that individual agents change their basket of goods according to the change of the “global utility” $U(\{\mathbf{x}^\alpha\})$ (i.e. a Hamiltonian) of the population rather than of their own utility. In other words, agents also take into account the change of utility of others when they update their choices. The main motivation behind this specification is that the dynamics will then spontaneously reach the Gibbs-Boltzmann equilibrium measure, Eq. (4.10). However, in the absence of a social planner coordinating all agents, such a dynamics is not very realistic. Yet, following an agent-based framework where individuals set out to improve their own utility function, including the herding component $c(\bar{x})^k$, would likely bring us in the realm of nonrelaxational dynamics, for which general analytical tools are still lacking.

Indeed, suppose we now consider dynamics where agents follow the logit rule but this time on their individual utility and not on the sum of individual utilities $U = \sum_\alpha u^\alpha$. The individual utility change of a randomly selected agent γ following a change in their basket of goods $\mathbf{x}^\gamma \rightarrow \mathbf{x}^\gamma + \Delta\mathbf{x}$ is

$$\Delta u^\gamma = \sum_i \frac{a_i}{x_i^\gamma} \left[1 + c(\bar{x}_i)^k + \frac{kc}{N} (\bar{x}_i)^{k-1} x_i^\gamma \log x_i^\gamma \right] \Delta x_i + O(\Delta\mathbf{x}^2). \quad (4.36)$$

Now, we assume that there exists a global utility function $\tilde{U}$ acting as the energy of the interacting system, meaning that at each timestep and for all agents we require

$$\begin{aligned} \Delta u^\gamma &= \tilde{U}(\mathbf{x}^1, \dots, \mathbf{x}^\gamma + \Delta\mathbf{x}, \dots, \mathbf{x}^N) - \tilde{U}(\mathbf{x}^1, \dots, \mathbf{x}^\gamma, \dots, \mathbf{x}^N) \\ &= \sum_{i=1}^M \frac{\partial \tilde{U}}{\partial x_i^\gamma} \Delta x_i + O(\Delta\mathbf{x}^2) \quad \forall \gamma \in [1, N], \end{aligned} \quad (4.37)$$

which by identification with Eq. (4.36) would give

$$\frac{\partial \tilde{U}}{\partial x_i^\gamma} = \frac{a_i}{x_i^\gamma} \left[1 + c(\bar{x}_i)^k + \frac{kc}{N} (\bar{x}_i)^{k-1} x_i^\gamma \log x_i^\gamma \right]. \quad (4.38)$$

Chapter 4. Slutsky matrices and the necessity of a global utility

An obvious requirement for $\tilde{U}$ (as for any other function) is that its Hessian must be symmetric. Taking the derivative with respect to another agent's consumption of the same product, it immediately becomes apparent that the required symmetry is violated due to the $\log x_i^\gamma$ term and that there can be no such function $\tilde{U}$. This contradiction therefore demonstrates that having individualistic agents maximize their own utility here means the system is non-Hamiltonian, and its dynamics are therefore nonrelaxational and necessarily detailed balance violating.

Natural questions are now (i) how may the current (equilibrium) model still be interpreted at the agent level and (ii) how is the phenomenology of the system impacted if we abandon the Hamiltonian description and take agents to maximize their own utilities.

To answer the former question, we start by writing the change in global utility ΔU if the randomly selected agent γ in the Monte Carlo dynamics accepts the proposed basket of goods $\mathbf{x}^\gamma + \Delta \mathbf{x}$. Assuming $N \gg 1$, we have

$$\Delta U \approx \sum_{i=1}^M \frac{a_i}{x_i^\gamma} \left(1 + c \left[1 + k \frac{x_i^\gamma}{\bar{x}_i} \overline{\log x_i} \right] (\bar{x}_i)^k \right) \Delta x_i, \quad (4.39)$$

where we remind the reader that the overline notation refers to arithmetic averages over the agents. By analogy with Eq. (4.37), this equation shows that the dynamics can in fact be interpreted in terms of agents only concerned with their own utility, albeit with modified values of the interaction parameter c ($c \rightarrow c_i^\gamma = c \left[1 + k \frac{x_i^\gamma}{\bar{x}_i} \overline{\log x_i} \right]$). The effective individual utility function is therefore configuration dependent and varies in between agents due to the difference in their consumption baskets, as somewhat expected given agents all coordinate to improve the overall outcome. Interestingly, the fact that agents maximize the global utility promotes the concentration phenomenon, as the logarithmic term will strongly penalize goods that have lost the favor of the crowd.

Regarding point (ii), simulating the system with a decision rule based on the individual utility but with a constant c – i.e. placing ourselves in a non-Hamiltonian setting – results in a largely unchanged phenomenology, with the same herding transition as was observed in the Hamiltonian case. The absence of this logarithmic penalty for small $\bar{x}_i$ appears to push the transition to slightly higher values of c and leads to more volatile Monte Carlo trajectories, with the appearance of more frequent switches in the vicinity of the transition. Although the absence of a solid theoretical framework to describe the steady-state in that case prevents us from drawing definitive conclusions at this stage, we conjecture that most of the results obtained at equilibrium regarding the Slutsky matrix continue to hold, with only minor quantitative modifications.

Alternatively, one could also purposefully write a global utility for which the detailed balance condition matches the maximization of an agent-specific utility.

To do so, the interaction term must be entirely symmetric in the sense that the change in global utility is identical regardless of the randomly selected agent (which was not the case with the previously studied model). For instance, one could take

$$U(\{\mathbf{x}^\alpha\}) = \sum_{i,\alpha} a_i^\alpha \log x_i^\alpha + \frac{1}{2} \sum_{\substack{i,\alpha,\gamma \\ \gamma \neq \alpha}} J_i^{\alpha\gamma} (x_i^\alpha)^\rho (x_i^\gamma)^\rho, \quad (4.40)$$

with $\mathbf{J}_i$ a *symmetric* interaction matrix and $0 < \rho < 1$ to preserve the concavity of the utility. In this case, we then have

$$\frac{\partial U}{\partial x_i^\alpha} = \frac{a_i}{x_i^\alpha} \left[1 + \rho (x_i^\alpha)^{\rho-1} \sum_{\substack{i,\gamma \\ \gamma \neq \alpha}} J_i^{\alpha\gamma} (x_i^\gamma)^\rho \right], \quad (4.41)$$

which is equal to the “selfish” derivative of utility of agent α defined as

$$u^\alpha(\{\mathbf{x}^\alpha\}) = \sum_i a_i \log x_i^\alpha + \sum_{\substack{i,\gamma \\ \gamma \neq \alpha}} J_i^{\alpha\gamma} (x_i^\alpha)^\rho (x_i^\gamma)^\rho. \quad (4.42)$$

(Note the factor $\frac{1}{2}$ difference between the definition of U and that of u^α .) This means that in such a model, the decision-making process based on a purely individualistic change of utility leads to an equilibrium distribution given by the Boltzmann weight $\exp(\beta U)$ (see [14] for a similar discussion in the context of the Schelling model mentioned in the introduction chapter).

The mean-field approximation of this model $J_i^{\alpha\gamma} = a_i J/N$ is studied in the limit $\beta \rightarrow \infty$ in App. B.4. For $\rho > \frac{1}{2}$, concentration will occur for sufficiently large values of J , and we therefore expect our results for the Slutsky matrix properties to remain largely unchanged.

In the case where $\mathbf{J}_i$ is not symmetric, the problem may no longer be treated with the standard techniques of equilibrium statistical mechanics. The system can then still be simulated with a decision rule maximizing the individual utility of the agent, but we can no longer assume that the system reaches an equilibrium given by the Gibbs-Boltzmann distribution Eq. (4.10). The impact of such non-reciprocal interactions will be at the heart of Part III below.

4.5.2 Equivalence of ensembles

In order to study analytically the equilibrium properties of our interacting set of agents and to determine when the herding transition occurs, we had to relax the budget constraint and place ourselves in the “grand-canonical” ensemble. As previously mentioned, the equivalence of results between the canonical and grand-canonical ensembles is not guaranteed *a priori*, although in the present case the




Figure 4.5: Evolution of the realized budget fluctuations in the grand-canonical ensemble as a function of c/c_{crit} for $M = 4$, $k = 2$, $w = 10$, $a_i = p_i = 1$ and $\beta = 0.5, 1, 2, 4, 10$, dotted lines showing the unconcentrated solution. Inset: log-log scale showing $\sigma^2 \beta$ decreasing as $1/c$ (dashed line) for $c \gtrsim c_{\text{crit}}$.

analytical (grand-canonical) results appear to match (canonical) simulations extremely well.

To formally assess the possible differences between the two ensembles, the budget fluctuations in the grand-canonical ensemble must be studied. The ensemble equivalence corresponds to cases where the variance of the realized budget are vanishingly small. Hence, we set out to compute

$$\sigma^2 = \left\langle \left(\sum_{i=1}^M x_i^\alpha p_i - w \right)^2 \right\rangle \quad (4.43)$$

in the grand-canonical ensemble. This computation may be performed by introducing small heterogeneous perturbations to the chemical potential, $\mu \rightarrow \mu + \delta\mu^\alpha$ and differentiating $\log \mathcal{Z}_N$ with respect to $\delta\mu^\alpha$. The calculation, detailed in App. B.3.4, finally yields

$$\sigma^2 = \sum_{i=1}^M \frac{(\bar{x}_i^* p_i)^2}{1 + \beta a_i [1 + c(\bar{x}_i^*)^k]}. \quad (4.44)$$

The evolution of this quantity for different values of β as a function of c in the previously discussed case $a_i = p_i = 1$ is shown in Fig. 4.5. As expected, this quantity vanishes for $\beta \rightarrow \infty$, as well as for $c \rightarrow \infty$. In the $c < c_{\text{crit}}$ region, where fluctuations are expected to be the strongest and thus where the difference between the ensembles is expected to be the largest, the non-condensed solution (dotted lines on Fig. 4.5) gives $\sigma^2 \sim O(M^{-1})$. As such, the grand-canonical and canonical ensembles will become strictly equivalent only in the limit $M \rightarrow \infty$. This being said, the Hessian of the free energy being independent of the chemical potential suggests that despite this apparent absence of strict equivalence between the ensembles for finite M , the onset of the transition appears to be

largely unaffected. In any case, although precise measures of the fluctuations of the system in the grand-canonical ensemble should be inaccurate for smaller values of β , this difference will not change the key results presented in this work. Besides, one should expect that some amount of budget fluctuations are indeed present in the real world!

4.6 Conclusion

Let us summarize what we have achieved in this chapter. By introducing a rationality parameter (or “intensity of choice”) β to account for the fact that agents are not strict utility maximizers, we have first reformulated the Slutsky equation within a general “fluctuation-response” framework, which allows one to express the Slutsky matrix in terms of consumption fluctuations only, without having to measure changes of consumptions when prices are slightly modified.

We have then shown that irrationality does *not* necessarily result in the breakdown of the symmetry of the Slutsky matrix. As a result, the hypothetical symmetry of empirically measured Slutsky matrices *cannot* be used as a general argument against bounded rationality [152].

When accounting for herding within large assemblies of agents, we found that symmetry is no longer guaranteed in general. Introducing a simple model of utility with interactions, we have indeed shown using the powerful methods offered by statistical mechanics that a concentration transition may occur, at which point strong selection of goods occurs. At the transition, the individual Slutsky matrix becomes markedly asymmetric, although the Slutsky matrix constructed using aggregate consumption can still remain symmetric when all agents are identical. Hence our result is not necessarily incompatible with existing work on non-unitary households [165]. From simulations, we also found that asymmetry is further amplified by out-of-equilibrium effects near the critical point. In line with standard consumer choice theory and most empirical studies, our model preserves the negative semi-definiteness of the Slutsky matrix regardless of interactions. Nonetheless and although not studied here, it should be noted that introducing some interactions in between products (representing redundancy or complementarity for example) in a similar framework appears to lead to positive eigenvalues [166].

Of course, further empirical studies on the properties of the Slutsky matrix would be of great interest, in particular to contrast our model of interacting, bounded rational agents with the recent sparsity-based approach of Gabaix [53]. To this end, we believe that fluctuation-response relations such as those presented in this work could be a very valuable econometric tool. Given the difficulty of conducting repeatable experiments in socioeconomic systems, the ability to estimate derivative quantities *without* actually requiring prices to change appears


Chapter 4. Slutsky matrices and the necessity of a global utility

quite promising.

Importantly in the context of this thesis, this first encounter with a socioeconomic model of interacting agents has also illustrated the limitations of equilibrium statistical mechanics. Indeed, in order to study the concentration transition and its effects on the Slutsky matrix analytically, we have implicitly assumed that agents are *altruistic* and all work at improving a *global* utility. While this is usually the rule in physics, where the energy is necessarily common to all constituents of the system, it is clear that a realistic description of human behavior should allow for individualistic dynamics. In Sec. 4.5.1, we have seen that the very existence of a global utility acting as the energy of the interacting system is the exception rather than the rule when the dynamics are prescribed at the agent level. This specificity of socioeconomic systems will be at the heart of the upcoming “SK-game” and of our version of the Schelling model in Chap. 7.

Key takeaways

- The Slutsky matrix describing the change in the demand of a product following the price change of another is known to be symmetric when agents are rational and noninteracting.
- Contrary to the common belief in economics, this symmetry is *not* a proof of rationality: the Slutsky matrix is also symmetric when modeling choices with the boundedly rational logit rule.
- Strong interactions, however, may lead to a violation of the symmetry of the matrix, particularly near a collective phase transition.



- While equilibrium statistical mechanics is appropriate to describe a single agent following the logit rule, this is often no longer the case if interacting agents are individualistic.

Chapter 4. Slutsky matrices and the necessity of a global utility

Part III

A unifying disordered model: the SK-game

Chapter 5

Presentation of the model and numerical results

All in the Game, yo... All in the Game.

Omar Little

So far, we have established that assuming agents to be purely rational is likely unrealistic due to the overwhelming complexity and associated uncertainty that may arise in even simple optimization problems. With this in mind, we have explored the consequences of having a probabilistic decision rule on simple non-interacting and interacting consumer choice models. Importantly, when agents are interacting, we have seen that describing the system with an equilibrium statistical mechanics formalism brings strong limitations to the possible dynamics that can be chosen at the agent level. Interactions between agents must notably be entirely reciprocal and constructed in a very particular way for a global utility function to be the suitable object to analyze the system's steady-state.

To go beyond these limitations we will now introduce a simple, unifying model, built at the agent level, and including not only bounded rationality and non-reciprocity but also *learning*.

5.1 Motivation

As discussed in the introduction chapter, classical economics is based on the idea that rational agents make optimal decisions, i.e. optimize their expected utility over future states of the world, weighted by their objective probabilities. Such an idealization of human behavior has been criticized by many (see e.g. [63, 167–170]). In particular, assuming that all agents are rational, allowing one to use game theoretic arguments to build such optimal strategies – often the result of complicated

mathematical calculations – is implausible, to say the least (see above).

A way to possibly save the rational expectation paradigm is to posit that agents are able to learn best responses from past experience. Yes, agents are only partially (“boundedly”) rational, but they learn and in the long run, they act “as if” they were rational [171]. This is clearly expressed by Evans and Honkapohja in their review paper on the subject [172]. They note that “[i]n standard macroeconomic models rational expectations can emerge in the long run, provided the agents’ environment remains stationary for a sufficiently long period.”

While seemingly reasonable, this proposition is by no means guaranteed to be legitimate. Indeed, the hypothesis that the environment should be stationary over “sufficiently long periods” can be restated in terms of the speed of convergence of the learning process, that should be short enough compared to the correlation time scale of the environment. However, in many circumstances and in particular in complex games, the convergence of the learning process to a collectively optimal state can be exceedingly long, or may in fact *never* take place. For example, reasonable learning rules can trap the system in some sub-optimal regions of the (high dimensional) solution space, see e.g. [173–176]. In other words, the learning process itself can be non-ergodic, even if the environment is described by an ergodic, stationary process. Another possibility is that agents’ strategies, even probabilistic, evolve chaotically forever, as was found by T. Galla & D. Farmer [177] in the context of competitive multi-choice two-player games, or by Bouchaud & R. Farmer [178] in a simple binary choice, multiplayer game. In such cases, the probabilities governing the different possible choices are not fixed but must themselves be described by probabilities.

This is in fact a generic feature of “complex systems”. As proposed by G. Parisi [179,180], the description of such systems requires the introduction of *probabilities of probabilities*, as their statistical behavior themselves (and not only individual trajectories) are highly sensitive to the small changes in parameters, initial conditions, or time. The inability to describe such systems with knowable probabilities was coined “radical complexity” in [181], a concept introduced in Chap. 3 above.

The sensitivity of optimal solutions to the parameters of the problem, or to the algorithm used to find them, has a very real consequence: one can no longer assume that all agents, even fully rational, will make the same decision, since any small perturbation may lead to a completely different solution, although similar in performance. In other words, the common knowledge assumption is not warranted. This has already been underlined in the context of portfolio optimisation in Chap. 3 and [42], or in the context of networked economies [176,182], but is expected to be of much more general scope, as anticipated by Keynes long ago and emphasized by many heterodox economists in the more recent past [63,168–170].

Here we want to dwell on this issue in the context of a multi-player binary

game – the “SK-game” –, understood as an idealization of the economic world where agents strongly interact in such a way that their payoffs depend on the action of others. In our setting, some relationships are mutually beneficial, while others are competitive. Agents have to learn how to coordinate to optimize their expected gains, which they do in a standard reinforcement way by observing the payoff of their actions and adapting their strategies accordingly.

5.2 A simple model for a complex world

5.2.1 Set-up of the model

As a minimal, stylized model for decision making in a complex environment of interacting agents, we restrict ourselves to binary decisions, as in many papers on models with social interactions, see e.g. [15, 58, 158, 178]. At every timestep t , each agent i plays $S_i(t) = \pm 1$, with $i = 1, \dots, N$, which can be thought of, for example, as the decision of an investor to buy or to sell the stock market, or the decision of a firm to increase or to decrease production, etc. The incentive to play $S_i(t) = +1$ is $Q_i(t)$ and is the agent’s estimate of the payoff associated to $S_i(t) = +1$ compared to that of $S_i(t) = -1$. The actual decision of agent i is probabilistic and drawn using the extensively discussed “logit” rule [55], i.e. sampled from a Boltzmann distribution over the choices of an agent,

$$\mathbb{P}[S_i(t) = \pm 1] = \frac{e^{\pm \beta Q_i(t)}}{e^{\beta Q_i(t)} + e^{-\beta Q_i(t)}} = \frac{1}{2}[1 \pm \tanh(\beta Q_i(t))], \quad (5.1)$$

or, equivalently, the expected choice (or “intention”) of agent i at time t is given by

$$m_i(t) := \langle S_i(t) \rangle = \tanh(\beta Q_i(t)). \quad (5.2)$$

Parameter β , assumed to be independent of i henceforth, is again analogous to the inverse temperature in statistical physics and represents the agent’s *rationality* or *intensity of choice*. As mentioned in Chap. 1 and in the previous chapter, the limit $\beta \rightarrow \infty$ corresponds to perfectly rational agents, that will systematically pick the choice that has their preference (given by the sign of $Q_i(t)$), while setting $\beta = 0$ gives erratic agents that randomly pick either decision with probability $1/2$ regardless of the value of $Q_i(t)$.

The evolution of the preference $Q_i(t)$ is where the learning takes place. We resort in so-called “ Q -learning” [183], i.e. reinforcement learning with a memory loss parameter α . Given the (yet unspecified) reward $\pm R_i(t)$ associated to making the choice ± 1 at time t , the evolution of incentives (and, in turn, beliefs) is given by

$$Q_i(t+1) = (1 - \alpha)Q_i(t) + \alpha R_i(t). \quad (5.3)$$

Chapter 5. Presentation of the model and numerical results

This map amounts to calculating an Exponentially Weighted Moving Average (EWMA) on the history of rewards $R_i(t)$. Taking $\alpha = 0$, the agent's preferences are fixed at their initial values, and we thus restrict ourselves to $\alpha > 0$. When $\alpha \rightarrow 0$, $Q_i(t)$ is approximately given by the average reward over the last α^{-1} time steps. Note here that this *averaging* of past rewards is not exactly the same as the *accumulation* rule (where the reward would not be multiplied by α in Eq. (5.3)) appearing in some forms of “Experience Weighted Attraction” that are popular in the socioeconomic context [184].

Now, the missing ingredient is the specification of the rewards, that encodes heterogeneity and non-reciprocity of interactions. Inspired by the theory of spin-glasses, in particular by the Sherrington-Kirkpatrick (SK) model of Sec. 2.2, we set

$$R_i(t) = \sum_{j=1}^N J_{ij} S_j(t). \quad (5.4)$$

Here, the matrix elements J_{ij} specify the mutually beneficial or competitive nature of the interactions between i and j . (Note that J_{ij} measures the impact of the decision of j on the reward of i .)

In the context of firm networks, a client-supplier relation would correspond to $J_{ij} > 0$, whereas two firms i, j competing for the same clients would correspond to $J_{ij} < 0$. In the so-called “Dean problem”, $J_{ij} > 0$ means that agents i and j get along well whereas $J_{ij} < 0$ means that they are in conflict [185]. The sign of S_i determines in which of the two available rooms agent i should sit, in order to minimize the number of possible conflicts. A predator-prey situation is when $J_{ij} \times J_{ji} < 0$, meaning that if i makes a gain, j makes a loss and vice versa. Importantly, whenever the interactions are non-symmetric, i.e. $J_{ji} \neq J_{ij}$, there cannot be a global utility function in the sense of Sec. 4.5.1 of the previous chapter. In other words, such an explicit non-reciprocity means the dynamics are nonrelaxational and the system non-Hamiltonian.

Note that the reward $R_i(t)$ depends on the actual (realized) decision of other players, and not their expected decisions or intentions. In other words, agents resort in *online* learning, which differs from *offline* learning where other players' decisions $S_i(t)$ are averaged over large batch sizes during which their inclinations would be assumed constant and replaced by their expectation $m_i(t)$.

Based on the learning dynamics, agents thus make a decision based on an imperfectly learned approximation of what other players are likely to do. Bringing equations (5.2), (5.3), (5.4) together, the evolution of agent i 's intention can be written as

$$m_i(t+1) = \tanh \left(\beta \sum_{j=1}^N J_{ij} \tilde{m}_j^\alpha(t) \right) \quad (5.5)$$

where $\tilde{m}_i^\alpha(t)$ is the estimate of agent j 's expected decision at time $t + 1$ based on its past actions up to time t . Explicitly writing the accumulated reward (EWMA) one has

$$\tilde{m}_i^\alpha(t) = \alpha \sum_{t' \leq t} (1 - \alpha)^{t-t'} S_i(t'). \quad (5.6)$$

Expressed in this form, it is clear that there is characteristic timescale $\tau_\alpha \sim 1/\alpha$ over which past choices contribute to the moving average $\tilde{m}_i^\alpha(t)$. Note that *offline* learning would correspond to a different evolution equation, namely

$$m_i(t + 1) = \tanh \left(\beta \sum_{j=1}^N J_{ij} m_j(t) \right), \quad (5.7)$$

although the two coincide in the $\alpha \rightarrow 0$ limit.

At this stage, it may be useful to compare and contrast the present model with previous work. On the one hand, the learning procedure closely resembles the original proposition by Sato & Crutchfield [186], and its treatment by Galla [187, 188] and others [189–191], however these authors considered games comprising only two players with many strategies. Subsequent cases explored by Galla and Farmer considering a larger number of players [177, 192] therefore lie closer to our setting but still consider many strategies, while the similar binary decision models proposed by Semeshenko *et al.* [193, 194] are restricted to perfectly rational agents and homogeneous interactions $J_{ij} = J > 0 \forall i, j$. Importantly, all these works also consider accumulated rewards and offline learning, in contrast with our averaged rewards and online learning. On the other hand, replicator models with random non-symmetric interactions between a large number of species [195, 196] share many features with the system at hand, but the prescribed dynamics are inherently linked to evolutionary principles such as extinction that are not present in our model. Finally, other Ising-inspired games such as that introduced in [197] are conceptually similar, in particular in their extension with myopic strategy revision (meaning updating based on future expectations and not directly passed realizations as done here) [198]. So far, however, these models have been studied without heterogeneities and therefore do not present the radical complexity related to the presence of a very large number of possible solutions discussed hereafter.

5.2.2 The interaction matrix

In order to rely on known results about the SK model, we will assume in the following that all agents randomly interact with one another, meaning that all elements of the matrix $\mathbf{J}$ are non-zero. Sparse matrices, corresponding to low-connectivity interaction matrices, would probably be more realistic in an economic

Chapter 5. Presentation of the model and numerical results

context. However, we expect that many of the conclusions reached below will qualitatively hold in such cases as well.

We choose interactions J_{ij} between i and j to be random Gaussian variables of order $N^{-1/2}$, with J_{ij} in general different from J_{ji} , accounting for possible non-reciprocity of interactions. More precisely, we introduce the parameter ε and write the interaction matrix as

$$J_{ij} = \left(1 - \frac{\varepsilon}{2}\right) J_{ij}^S + \frac{\varepsilon}{2} J_{ij}^A, \quad (5.8)$$

with $\mathbf{J}^S$ a symmetric matrix and $\mathbf{J}^A$ an anti-symmetric matrix. The entries of both these matrices independent and sampled from a Gaussian distribution of mean 0 and variance σ^2/N . This defines what we will call the “SK-game” henceforth. A non-zero average value of $\mathbf{J}^S$ will be discussed below.

In the following we set $\sigma = 1$ without loss of generality. The resulting variance of J_{ij} is thus given by

$$v(\varepsilon) := N \text{Var}(J_{ij}) = 1 - \varepsilon + \frac{1}{2}\varepsilon^2. \quad (5.9)$$

The specific cases $\varepsilon = \{0, 1, 2\}$ hence correspond to fully symmetric ($J_{ij} = J_{ji}$), a-symmetric (i.e. J_{ij} and J_{ji} independent) and anti-symmetric ($J_{ij} = -J_{ji}$) interactions respectively. We can thus also characterize the correlation between J_{ij} and J_{ji} through parameter η ,

$$\eta = \frac{\overline{J_{ij}J_{ji}}}{\overline{J_{ij}^2}} = \frac{1 - \varepsilon}{v(\varepsilon)}, \quad (5.10)$$

where overlines indicate an average over the disorder.

It may actually be insightful to allow for a non-zero average value to the interaction parameters, and define the reward $R_i(t)$ as

$$R_i(t) = \sum_{j=1}^N J_{ij} S_j(t) + J_0 M(t); \quad M(t) := \frac{1}{N} \sum_{j=1}^N S_j(t). \quad (5.11)$$

Finally, one may think that agents have some idiosyncratic preferences, or different costs associated to the two possible decisions $S_i = \pm 1$. This would amount to adding to the reward $R_i(t)$ a time independent term H_i , where H_i favors $S_i = +1$ if positive and $S_i = -1$ if negative, see the RFIM discussed in Sec. 1.3.2. Here, we will restrict to $H_i \equiv 0, \forall i$, but one expects from the literature on spin-glasses that the main results discussed below would still hold for small enough H_i s. Beyond some threshold value, on the other hand, agents end up aligning to their *a priori* preference, i.e. $m_i H_i > 0$.

5.3 Overview and numerical results

Having established the rules of the SK-game, a first and natural step is to simulate the evolution of the system from $Q_i(0) = 0 \forall i$ at finite values of N . Exploring the parameter space in this way, we can establish the broad types of behaviors displayed by the model. In the following, we will summarize these results before jumping into a more technical analysis in the following chapter.

5.3.1 A word on the average reward

In the SK-game, the payoff of each agent is a random function of the decisions of all other agents. Hence, learning the optimal strategy (in terms of the probability for agent i to play $+1$ or -1) is bound to be extremely difficult.

Defining the average reward as

$$\mathcal{R}_N := \frac{1}{N} \sum_i S_i R_i = \frac{1}{N} \sum_{i,j=1}^N J_{ij} S_i S_j, \quad (5.12)$$

and noting that $\sum_{i,j=1}^N J_{ij} S_i S_j = (1 - \frac{\varepsilon}{2}) \sum_{i,j=1}^N J_{ij}^S S_i S_j$, the largest possible average reward $\mathcal{R}_\infty$ for $N \rightarrow \infty$ can be exactly computed using the celebrated Parisi replica-symmetry breaking scheme of the classical SK model and reads [44]:

$$\lim_{N \rightarrow \infty} \mathcal{R}_N := \mathcal{R}_\infty = 0.7631... \times (2 - \varepsilon). \quad (5.13)$$

However, in practice there is no known algorithm to find the global minimum for a given draw of the disorder in polynomial time in N – although it has been shown that one can devise an algorithm with time growing like $K(\epsilon)N^2$ to find configurations $\{S_i\}$ that reach a value of at least $1 - \epsilon$ the optimum (5.13), $\epsilon > 0$ [199]. In any case, it is expected that simple learning algorithms will inevitably fail to find the true optimal solution. Nevertheless, we also know from the spin-glass folklore (see below for more precise statements) that many configurations of $\{S_i\}$'s correspond to quasi-optima, or, in the language of H. Simon, satisficing solutions [63]. It is in a sense the proliferation of such sub-optimal solutions that prevent simple algorithms to find the optimum optimorum. Furthermore, if learning indeed converges (which is not the case when ε is too large, i.e. when interactions are not reciprocal enough), the obtained fixed point heavily depends on the initial condition and/or the specific interaction matrix $\mathbf{J}$.

5.3.2 Phase diagram in the noiseless limit

Let us first consider the case where agents always choose the action that would have had the best average reward in the past $Q_i(t)$ (this assumes that other agents




Figure 5.1: Qualitative phase diagram in the (ε, α) plane for the SK-game in the (a) noiseless ($\beta \rightarrow \infty$) and (b) weak noise ($\beta \gg 1$) regimes. FP refers to “fixed point”, “Para” means that players’ strategies are completely random, whereas “chaos” means that at a given instant of time players have well defined intentions m_i but these evolve chaotically with time.

still played what they played). This corresponds to the noiseless learning limit $\beta \rightarrow \infty$. In this case, the iteration map Eq. (5.5) becomes

$$S_i(t+1) = \text{sign} \left(\sum_{t' \leq t} (1-\alpha)^{t-t'} \sum_{j=1}^N J_{ij} S_j(t') \right), \quad (5.14)$$

and the model is fully specified by two parameters: α (controlling the memory time scale of the agents) and ε (controlling the reciprocity of interactions). For N not too large, the evolution of Eq. (5.14) leads to either fixed points, or oscillations, or else chaos. The schematic phase diagram in the plane (α, ε) is shown in Fig. 5.1(a).

One clearly sees a region for (α, ε) small where learning reaches a fixed point, where the average reward is close, but significantly below the theoretical optimum $\mathcal{R}_\infty$ given by Eq. (5.13), see Fig. 5.2.²⁵ Note that learning definitely helps: for $\varepsilon = 0$, most fixed points are characterized by a typical reward $\mathcal{R} \approx 1.01$ (which can be computed by taking an unweighted average on all TAP states, see Sec. 2.2.2) [106], significantly worse than the value ≈ 1.40 reached by our learning agents extrapolated to $N \rightarrow \infty$.

As ε and α are varied one observes the following features:

- When ε is not too large (interactions sufficiently reciprocal) and α increases (shorter and shorter memory) learning progressively ceases to converge and oscillations start appearing: impatient learning leads to cycles. This leads

²⁵Note that there are finite N corrections that must be taken into account for such a comparison, which read [200] $\mathcal{R}_N \approx \mathcal{R}_\infty - AN^{-2/3}$, $A \approx 0.75 \times (2 - \varepsilon)$, see also Fig. 5.4(b).

to a sharp decrease of the average reward (see Fig. 5.2(a)), as agents over-react to new information and are no longer able to coordinate on a mutually beneficial equilibrium. A similar effect was observed in a dynamical model of firm network, where over-reaction leads to oscillating prices and production [150] (see also [182]).

- Conversely, when α is small (long memory) and ε increases, the probability to reach a fixed point progressively decreases, and when a fixed point is reached, the average reward is reduced see Fig. 5.2(b). Beyond some threshold value, the dynamics becomes completely chaotic, leading to further loss of reward. Note that “chaos” here means that although at any given instant of time, agents have well defined intentions $m_i(t) \neq 0$, these intentions evolve chaotically forever.²⁶
- Surprisingly, oscillations *reappear* when ε becomes larger than unity, i.e. when interactions are mostly anti-symmetric, “predator-prey” like. Perhaps reminiscent of the famous Lotka-Volterra model, agents’ decisions and pay-offs become periodic, with a period that scales anomalously as $\alpha^{-1/2}$, i.e. much shorter than the natural memory time scale α^{-1} (see Chap. 6 below). Although not the Nash equilibrium²⁷ $m_i \equiv 0$, these oscillations allow the average reward to be positive when ε is not too large, although at each instant of time, some agents have negative rewards. In these cases, the system therefore self-organizes in such a way that the *collective* outcome outperforms the Nash equilibrium, despite the agents being individualistic.
- Only close to the extreme competition limit $\varepsilon = 2$ (corresponding to a zero sum game, cf. Eq. (5.13)) and for small α , are agents able to learn that the unique Nash equilibrium is to play random strategies $m_i \equiv 0$ (see Fig. 5.1, bottom right region).
- Finally, in the extreme (and unrealistic) case $\alpha = 1$, where agents choose their strategy based on the last reward only, the system evolves, as ε increases from zero, from high frequency oscillations with period $L = 2$ to “weak chaos” to “strong chaos” when $\varepsilon \approx 1$ and finally back to oscillations of period $L = 4$ when $\varepsilon \rightarrow 2$.

In order to characterize more precisely such temporal behaviors, it is useful to (re)introduce the two-point auto-correlation function of the *expected* decisions or

²⁶Note that the role of ε is somewhat similar to that of parameter Γ in [177]: increasing competition leads to chaos.

²⁷The Nash equilibrium is the solution found such that no individual can improve their outcome by changing their choice, see [49].




Figure 5.2: Evolution of the average reward with (a) the memory loss rate α , for $\varepsilon = \{0, 0.6, 0.85, 1.05, 1.5, 2\}$ from dark purple to light green, $\beta \rightarrow \infty$; (b) the asymmetry ε for $\alpha = 0.01$ and $\alpha = 0.1$, $\beta \rightarrow \infty$, the vertical dotted line indicates the value ε_c for which the system becomes chaotic [201]. In both (a) and (b), $N = 256$ and the dashed line represents the Parisi solution for $\mathcal{R}_\infty$.

intentions:

$$C(t, t + \tau) = \frac{1}{N} \sum_i \overline{\langle m_i(t) m_i(t + \tau) \rangle}, \quad (5.15)$$

where the angular brackets now refer to an average over initial conditions.²⁸ In cases where the dynamics are assumed to be time-translation invariant, we will write $C(\tau)$ which corresponds to the above quantity averaged over time after the system has reached a steady-state.

The autocorrelation function corresponding to the different cases described above are plotted in Fig. 5.3. Note that the signature of oscillations of period L is that $C(nL) \equiv 1$ for all integer n . However, note that when $\varepsilon < \varepsilon_c$, not all spins flip at each time step. The fact that $C(2n + 1) = 0$ means that half of the spins in fact remain fixed in time, while the other half oscillate in sync between $+S_i$ and $-S_i$.²⁹ In the chaotic phases, $C(\tau)$ tends to zero for large τ , with either underdamped or overdamped oscillations. Hence in these cases, the configuration $\{S_i\}$ evolves indefinitely with time, and hardly ever revisit the same states.

5.3.3 Noisy learning

In the presence of noise, the “convictions” $|m_i|$ of agents naturally decrease, and in fact become zero (i.e. decisions are totally random) beyond a critical noise level that depends on the asymmetry parameter ε : more asymmetry leads to more

²⁸Of course, in the $\beta \rightarrow \infty$ limit discussed in this subsection, one can replace $m_i(t)$ by the actual decision $S_i(t)$ in the definition of $C(t, t + \tau)$. More generally, the spin-spin correlation function is given by $(1 - q(t))\delta(\tau) + C(t, t + \tau)$ with $q(t) = \overline{\langle m_i^2(t) \rangle}$.

²⁹Note the rather large error bar on $C(2n + 1)$, meaning that there are actually substantial fluctuations of the number of idle spins around the value $N/2$ when N is finite.




Figure 5.3: Evolution of the steady-state two-point correlation function for $\alpha = 1$, $\beta \rightarrow \infty$, markers indicating simulations of the game at $N = 256$ averaged over 128 samples of disorder and initial conditions, error-bars showing 95% confidence intervals. The continuous lines representing the solution to the Dynamical Mean-Field Equations that will be discussed in Chap. 6 below. (a) $\varepsilon = 0.1$ ($\varepsilon < \varepsilon_c \approx 0.8$), cycles of length $L = 2$. (b) $\varepsilon = 0.85$ ($\varepsilon_c < \varepsilon < 1$), “weakly” chaotic behavior. (c) $\varepsilon = 1.05$ ($1 < \varepsilon < 2 - \varepsilon_c$), “strongly” chaotic behavior. (d) $\varepsilon = 1.5$ ($\varepsilon > 2 - \varepsilon_c$), cycles of length $L = 4$.




Figure 5.4: Evolution of the average reward for $N = 256$ with (a) the noise level $1/\beta$ for $\alpha = 0.01$, $\varepsilon = \{0, 0.6\}$ (dark purple and blue respectively), note the non-monotonic behavior. The dashed line represents the Parisi solution for $\mathcal{R}_\infty$; (b) system size N for $\alpha = 0.01$, $\beta \rightarrow \infty$, $\varepsilon = \{0, 0.6\}$ (dark purple and blue respectively), continuous lines showing fits $\mathcal{R}_N = \mathcal{R}_\infty - AN^{-2/3}$, excluding $N = 32$. The dotted line indicates the best fit for the SK ground state, for which $A \approx 1.5$ [200].

fragile convictions (see Fig. 6.9 below for a more precise description). Because of the noise, strict fixed points do not exist anymore, but are replaced (for ε small enough) by quasi-fixed points – in the sense that the intentions m_i fluctuate around some plateau value for very long times, before evolving to another configuration completely uncorrelated with the previous one. This process goes on forever, albeit at a rate that slows down with time: plateaus become more and more permanent. This corresponds to the aging phenomenon introduced in Chap. 2 for the original SK model, and will be discussed in more detail in a dedicated section of the next chapter.

In a socio-economic context, it means that a form of quasi-equilibrium is temporarily reached by the learning process, but such a quasi-equilibrium will be completely disrupted after some time, even in the absence of any exogenous shocks. This is very similar to the quasi-nonergodic scenario recently proposed in Ref. [178], although in our case the evolution time is not constant but increases with the “age” of the system, i.e. the amount of time the game has been running.

Perhaps counter-intuitively, however, the role of noise is on average beneficial when $1/\beta$ is not too large. Indeed, as shown in Fig. 5.4(a), the average reward first increases as a small amount of noise is introduced, before reaching a maximum beyond which “irrationality” becomes detrimental. The intuition is that without noise, the system gets trapped by fixed points with large basins of attraction, but low average rewards. A small amount of noise allows agents to reassess their intentions and collectively reach more favorable quasi-fixed points, much as with simulated annealing.

When learning leads to a chaotic evolution, i.e. when J_{ij} and J_{ji} are close to uncorrelated ($\varepsilon \sim 1$), noise in the learning process does not radically change the evolution of the system: deterministic chaos just becomes noisy chaos. However, there is still a distinction between a low-noise phase where at each instant of time, agents have non-zero expected decisions m_i (that will evolve over time) from a high-noise phase where agents always make random choices between ± 1 with probability $1/2$ (see Fig. 5.1).

Finally, in the case where learning leads to cycles, any amount of noise irretrievably disrupts the synchronisation process, and cycles are replaced by pseudo-cycles, with either underdamped or overdamped characteristics. In the limit $\varepsilon \rightarrow 2$, any level of fluctuations drive the system to a paramagnetic state where $q = C(0) = 0$ (see Fig. 6.10), meaning the agents remain undecided (corresponding to the previously mentioned Nash equilibrium).

5.3.4 Distribution of individual rewards

As we have noted above, the average reward is close, but significantly below the theoretical optimum $\mathcal{R}_\infty$ given by Eq. (5.13). However, some agents are better




Figure 5.5: Distribution of individual rewards at $\beta \rightarrow \infty$ for $N = 512$ and 128 initial conditions and realizations of the disorder in the fully symmetric case $\varepsilon = 0$ for $\alpha = \{0.5, 0.1, 0.01\}$ (dark to light coloring). The dashed line is the Sommers-Dupont analytical solution to the SK model [202]. Inset: associated survival function in a lin-log scale and focusing on the right tail, dotted line corresponding to a Gaussian fit.

off than others, in the sense that the individual reward e_i at the fixed point (when fixed points exist), is different from agent to agent. Noting that

$$e_i := S_i^* R_i^* = \left| \sum_j J_{ij} S_j^* \right|, \quad (5.16)$$

where the second equality holds because at the fixed point one must have $S_i^* = \text{sign}(R_i^*)$, it is clear that in the fully reciprocal case $\varepsilon = 0$, all rewards e_i are positive.

The *distribution* $\rho(e)$ of these rewards over agents is expected to be self-averaging for large N , i.e. independent of the specific realization of the J_{ij} and of the initial condition. Such distribution is shown in Fig. 5.5. One notices that $\rho(e)$ vanishes linearly when $e \rightarrow 0$

$$\rho(e) \underset{e \rightarrow 0}{\approx} \kappa e, \quad \kappa \approx 1.6,$$

as for the standard SK model, although the value of κ is distinctly different from the one obtained for the true optimal states of the SK model, for which $\kappa_{\text{SK}} \approx 0.6$ [202, 203]. Such a discrepancy is expected, since the fixed points are obtained as the long time limit of the learning process – in particular, since $\kappa > \kappa_{\text{SK}}$, the number of poorly rewarded agents is too high compared to what it would be in an optimal state. Note that once α is sufficiently small for the system to reach a fixed point, its precise value does not seem to have an impact on the distribution of rewards and κ .

Another important remark is that the distribution of rewards $\rho(e)$ does not develop a “gap” for small e , i.e. a region where $\rho(e)$ is exactly zero. In other words,




Figure 5.6: Evolution of individual rewards in time for $N = 256$, $\alpha = 0.01$, $\beta \rightarrow \infty$, (a) $\varepsilon = 0$ and (b) $\varepsilon = 0.6$. Right: histogram of the individual rewards after a single timestep (shaded) and at the final time (unshaded).




Figure 5.7: Distribution of individual rewards for $N = 256$ and $\alpha = \{1, 0.5, 0.01\}$ represented by black triangles, purple squares and green circles respectively, $\beta \rightarrow \infty$, measured over 32 initial conditions and realizations of the disorder. From left to right: $\varepsilon = \{0, 0.1, 0.85, 1.05, 1.5, 2\}$.

although all agents have positive rewards, some of them are very small. This is associated with the so-called “marginal stability” of the equilibrium state [204], to wit, its fragility with respect to small perturbations, as discussed in more details in the next subsection.

For very large e , the distribution $\rho(e)$ decreases like a Gaussian (Fig. 5.5 inset), corresponding to a Central Limit Theorem behavior in that regime, as for the SK model.³⁰ Fig. 5.6 shows how the rewards of individual agents evolve from an initially random configuration, before settling to constant (but heterogeneous) values at the fixed point.

As competitive effects get stronger (i.e. as ε increases) and the system ceases to reach a fixed point, the distribution $\rho(e)$ develops a tail for negative values of e , meaning that some agents are systematic losers, see Fig. 5.7. In the extreme “predator-prey” limit $\varepsilon = 2$, the distribution $\rho(e)$ becomes perfectly symmetric around $e = 0$, as expected – see Fig. 5.7, rightmost plot.

³⁰One would expect a different behavior in the possibly relevant case of a fat-tailed distribution of the J_{ij} , see [205–207]. We leave this question for further investigations.




Figure 5.8: Overlap between solutions for different initial conditions and identical draws of interactions, for $\alpha = 0.01$, $\beta \rightarrow \infty$, $\varepsilon = 0$. (a) Distribution of overlaps shifted by the mean and rescaled with the system size N to the power $2/3$. (b) Average overlap as a function of system size in log-log coordinates, with the best regression line $N^{-0.85}$. Error-bars show the 95% confidence interval over 16 different draws of the disorder.

5.3.5 Unpredictability of equilibria

Now, the interesting point about our model is that the final rewards are highly dependent on the initial conditions and/or the realization of the J_{ij} 's. In other words, successful agents in one realization of the game become the losers for another realization obtained with different initial conditions. A way to quantify this is to measure the cross-sectional correlation of final rewards for two different realizations, i.e.

$$C_N^\times := \frac{1}{N} \sum_i (e_i^a - \langle e^a \rangle)(e_i^b - \langle e^b \rangle), \quad a \neq b \quad (5.17)$$

where a, b corresponds to two different initial conditions and $\langle e \rangle$ corresponds to the cross-sectional average reward. As shown in Fig. 5.8, $C_N^\times$ goes to zero at large N , indicating that the final outcome of the game, in terms of the winners and the losers, cannot be predicted. The dependence on N appears to be non-trivial, with different exponents governing the decay of the mean overlap $\overline{C_N^\times}$ (decaying as $N^{-0.85}$) and its standard deviation (decaying as $N^{-2/3}$).

A similar effect would be observed if instead of changing the initial condition one would randomly change the interaction matrix $\mathbf{J}$ by a tiny amount ϵ . The statement here is that for any small ϵ , $C_N^\times$ goes to zero for sufficiently large N . This is the “disorder chaos” discussed in the portfolio problem of Chap. 3; by analogy with known results for the SK model, we conjecture that $C_N^\times$ is a decreasing function of $N\epsilon^\zeta$, where $\zeta = 3$ in the SK case [208]. This means that when $N \gg \epsilon^{-\zeta}$ the rewards between two systems with nearly the same interaction structure, starting with the same initial conditions, will be close to independent.

Such a sensitive dependence of the whole steady-state of the system (in our case the full knowledge of the intentions m_i of all agents) once again prevents any kind of “common knowledge” assumption about what other agents will decide to do in a specific environment. No reasonable learning process can lead to a predictable outcome; even the presence of a benevolent social planner assigning their optimal strategy to all agents would not be able to do so without a perfect knowledge of all interactions between agents and without exponentially powerful (in N) computing abilities. Such a “radically complex” situation leads to “radical uncertainty” in the sense that the behavior of agents, even rational, cannot be predicted. Learning agents can only achieve satisficing solutions, that are furthermore hypersensitive to details. Moreover, as we have seen in Sec. 5.3.3, any amount of noise in the learning process will make the whole system “jump” from one satisficing solution to another in the course of time.

5.3.6 Increasing cooperativity

A way to help agents coordinate is to use rewards given by Eq. (5.11) with $J_0 > 0$, representing a non-zero average cooperative contribution to rewards. This term obviously helps agents finding mutually beneficial strategies. (Note that with our normalisation, the J_0 term is in fact $N^{-1/2}$ times smaller than the random interaction terms J_{ij} .)

The impact of such a term is well understood in the case of the SK model for $\varepsilon = 0$ [209]. For $\beta = \infty$, one finds that whenever $J_0 \leq 1$, the average intention $M(t)$ remains zero for large N and one expects that the learning process is not affected by such a “nudge”. When $J_0 > 1$, on the other hand, the situation changes as all agents start coordinate on one of the two possible choices: the average intention becomes non-zero, although a finite fraction of agents still play opposite to the majority because of their own idiosyncratic rewards.

For $J_0 \gg 1$, radical complexity disappears and learning quickly converges to the obvious optimal strategy where all agents make the same move $S_i = +1$ or $S_i = -1$, $\forall i$. In this case, $\mathcal{R}_N = J_0$ as $M(t)$ eventually reaches unity, see Fig. 5.9. For $\varepsilon > 0$, the same occurs barring some rescaling, as visible in Fig. 5.9.

In the case $J_0 < 0$ with $|J_0| \gg 1$, the only solution of Eq. (5.7) (valid for $\alpha \rightarrow 0$) is $m_i = 0$ for all i , i.e. agents cannot coordinate and play random strategies.

5.3.7 Self-reinforcement and habit formation

Up to this point, all results have assumed that there is no self-interaction, $J_{ii} = 0$. Nonetheless, it is interesting to consider the possibility of having an $O(1)$ positive diagonal term in the interaction matrix. In the socioeconomic context, such a contribution is relevant as it represents self-reinforcement of past choices, which




Figure 5.9: (a) Evolution of the average reward with the incentive to cooperate J_0 for $N = 512$, $\alpha = 0.01$, $\beta \rightarrow \infty$ and $\varepsilon = \{0, 0.6, 0.85, 1.05, 1.5, 2\}$ with colors ranging from purple to light green with increasing ε . (b) Average intention in the long time limit $M = \lim_{t \rightarrow \infty} M(t)$ for the same parameters. Note the rather strong finite size effects, in particular for $\varepsilon = 0$.

is also called “habit formation” where agents stick to past choices, a popular idea in behavioral science, see e.g. [173, 175] and refs. therein.

The introduction of a diagonal contribution has important consequences for the problem. Assuming the self-interaction is identical for all agents, $J_{ii} = J_d > 0$, it will rather intuitively favor the emergence of fixed points since agents will be tempted to stick to past choices. It is for instance known that in the case of fully random interactions $\varepsilon = 1$, fixed points will start to appear when J_d is sufficiently large [210]. Interestingly, these fixed points can be very difficult to reach dynamically with standard Hopfield dynamics ($\alpha = 1$).

Adding such diagonal contribution to our learning dynamics, we have observed that the fraction of trajectories converging to seemingly dynamically inaccessible configurations significantly increases, specially when $\alpha \ll 1$. While further work would be required to precisely assess the effectiveness of learning when self-reinforcement is present, particularly as finite size effects appear to play a significant role, such a result is consistent with the overall influence of learning reported here.

5.4 Core message

Before jumping into a more technical analysis of the different regimes of the SK-game, let us summarize the broad conclusions we draw in terms of socioeconomic modeling. In line with the core message of Galla & Farmer [177], our multi-agent binary decision model provides an explicit counter-example to the idea advocated in [172] that learning could save the rational expectation framework (cf. Sec. 5.1).

Learning in general does not converge to any fixed point, even when the

Chapter 5. Presentation of the model and numerical results

environment (in our case the interaction matrix $\mathbf{J}$) is completely static: non-stationarity is self-induced by the complexity of the game that agents are trying to learn, as also recently argued in [176].


When learning does indeed converge (which requires *a minima* a high level of reciprocity between agents) the collective state reached by the system is far from the optimal state, which only a benevolent, omniscient social planner with formidable powers can achieve. In other words, even more sophisticated learning rules would not really improve the outcome: the SK-game is *unlearnable* and – as argued by H. Simon [63] – agents must resort to *satisficing* solutions.

Furthermore, any small random perturbation (noise in the learning process, or slow evolution in the environment) eventually destabilizes any fixed point reached by the learning process, and completely reshuffles the collective state of the system: in the long run, agents initially favoring the + end up favoring –, and better-off agents end up being the underdogs, and vice-versa (much as in the simpler model of Ref. [178]).

Finally, even in the most favorable case of a fully reciprocal game with slow learning, the average reward is in fact *improved* when some level of noise (or irrationality) is introduced in the learning rule, before degrading again for large noise.

Key takeaways

- The SK-game is a unifying model of a *radically complex* world in which N players make a binary choice at every time step:
 - The incentive to play either ± 1 is based on past rewards, that agents forget with a *memory loss rate* α ,
 - Agents make their decision following the logit rule with a *rationality parameter* β ,
 - Rewards are based exclusively on the decisions of others and are non-reciprocal, with a payoff *asymmetry parameter* ε .
- When agents have long term memory and interactions are sufficiently reciprocal, intentions reach “satisficing” (and fragile) fixed points.
- The average reward in these fixed points is markedly below the best possible solution, but significantly above the memoryless outcome.
- When interactions are close to completely uncorrelated, the agents’ intentions evolve chaotically.
- When interactions are mostly non-reciprocal, intentions display oscillations that do not correspond to the Nash equilibrium.
- In the fixed point regime, irrationality is initially somewhat beneficial, but leads to permanent non-stationarity in the form of aging.



- Overall, learning cannot save the rational expectation framework when radical complexity is involved.

Chapter 5. Presentation of the model and numerical results

Chapter 6

Detailed analysis of the model

In order to solve this differential equation you look at it till a solution occurs to you.

George Polya

We have seen that the SK-game displays a wide variety of complex collective behavior. In this chapter, we delve into the detailed analysis of the model's statics and dynamics, and attempt to describe its rich phenomenology analytically.

6.1 Fixed point analysis and complexity

Only in some cases does learning converge to non-trivial fixed points where strategies are probabilistic but with time independent probability $p_{\pm}$ to play ± 1 , such that $p_{\pm} = (1 \pm m_i^*)/2$ for agent i . Such a steady-state would be analogous to an *economic equilibrium* (although it is essential to dissociate this notion from that of a thermodynamic equilibrium, which may only exist in the case of fully reciprocal interactions, $\varepsilon = 0$).

We will mostly focus, in the following, on the long term memory case $\alpha \ll 1$ which is most relevant for thinking about learning in a (semi-)realistic context. In this case, one can show that the exponential moving average on the realized values $S_i(t)$ converges to one on the expected values $m_i(t)$. Indeed, as detailed in Appendix C.1,

$$\left\langle \left(\alpha \sum_{t' \leq t} (1 - \alpha)^{t-t'} (m_i(t') - S_i(t')) \right)^2 \right\rangle \leq \frac{\alpha}{2 - \alpha} \xrightarrow{\alpha \rightarrow 0} 0. \quad (6.1)$$

This means that up to fluctuations of order $\sqrt{\alpha}$, we can describe the dynamics of the system through a deterministic iteration on $m_i(t)$. (In fact, we will see below that the neglected fluctuations are of order $\sqrt{\alpha/\beta}$.)

Chapter 6. Detailed analysis of the model

Further making the *ansatz* that the mean-field dynamics will eventually reach a fixed point $m_i(t) = m_i^* \forall i$ given sufficient time, Eq. (5.5) then yields

$$m_i^* = \tanh \left(\beta \sum_j J_{ij} m_j^* \right). \quad (6.2)$$

This is none other than the Naive Mean-Field Equations (NMFE) introduced in Sec. 2.2.2 when $\varepsilon = 0$, and defines a so-called static Quantal Response Equilibrium, similar to its fully mean-field equivalent ($J_{ij} = J/N$) studied in [197].

As previously mentioned, the NMFE are very similar to the celebrated TAP equations [103] describing the mean magnetization in the Sherrington-Kirkpatrick (SK) spin-glass [28]. Physically, the NMFE are satisfied when extremizing the free energy of a system of N sites comprising $M \rightarrow \infty$ binary spins, with sites interacting through an SK-like Hamiltonian [104]. Despite being seemingly simpler than its previously mentioned TAP counterpart, which includes an additional Onsager “reaction term”, the NMFE share many of its properties. Relevant to our problem, both the NMFE and the TAP equations have a paramagnetic phase ($m_i^* = 0 \forall i$) for $\beta < \beta_c$, while above this critical value there is a spin-glass phase where $q^* = N^{-1} \sum_i (m_i^*)^2 > 0$ and solutions are exponentially abundant in N [104, 211, 212]. Recall the NMFE has a critical temperature $1/\beta_c = 2$ as opposed to $1/\beta_c = 1$ in the TAP case, while the two equations become strictly equivalent in the $\beta \rightarrow \infty$ limit.

Using known properties from the spin-glass literature, we can therefore already establish that *if* the system reaches a fixed point when interactions are fully reciprocal ($\varepsilon = 0$) and memory is long ranged, it will be either a trivial fixed point where agents continue making random decisions for ever ($m_i^* = 0 \forall i$), or, when learning is not too noisy ($\beta > \beta_c$) the number of fixed points is $\sim \exp[\Sigma(\beta)N]$, where we remind $\Sigma(\beta)$ is called the “complexity”. In this second case, the fixed point actually reached by learning depends sensitively on the initial conditions and the interaction matrix $\mathbf{J}$.

How is this standard picture altered when interactions are no longer reciprocal? In such cases, the system cannot be described using the equilibrium statistical mechanics machinery.

6.1.1 Critical noise level

In order to extend the notion of critical noise β_c to $\varepsilon > 0$, one can naively look at the linear stability of the paramagnetic solution $m_i^* = 0 \forall i$ to Eq. (6.2). Just as in the TAP case [44], expanding the hyperbolic tangent to the second order and projecting the vector of m_i on an eigenvector of $\mathbf{J}$, the stability condition can be expressed with the largest eigenvalue of the interaction matrix. Adapting

known results from random matrix theory to our specific problem formulation, the spectrum of $\mathbf{J}$ can be expressed as an interpolation between a Wigner semi-circle on the real axis ($\varepsilon = 0$), the Ginibre ensemble ($\varepsilon = 1$) and a Wigner semi-circle on the imaginary axis ($\varepsilon = 2$) [213]. The resulting critical “temperature” is then given by

$$T_c(\varepsilon) = \frac{1}{\beta_c(\varepsilon)} = \frac{1}{2} \frac{(2 - \varepsilon)^2}{\sqrt{v(\varepsilon)}}, \quad (6.3)$$

recovering the known result $1/\beta_c = 2$ for the case $\varepsilon = 0$. (We recall that we have set the interaction variance σ^2 to unity. If needed σ can be reinstalled by the rescaling $\beta \rightarrow \beta\sigma$.)

6.1.2 The elusive complexity

To determine if there are still an exponential number of fixed point to reach below the candidate critical noise level, i.e. if there is a spin-glass phase, for $\beta > \beta_c(\varepsilon)$ when $\varepsilon > 0$, we should study the complexity, defined just as in the portfolio problem,

$$\Sigma(\beta, \varepsilon) = \lim_{N \rightarrow \infty} \frac{1}{N} \log \mathcal{N}_J(N, \beta, \varepsilon), \quad (6.4)$$

where $\mathcal{N}_J$ is the number of fixed points in the system for a given interaction matrix. As discussed above, there are then two ways to compute an average of this quantity over the disorder: the “quenched” complexity, where the mean of the logarithm of the number of solutions is considered, and its “annealed” counterpart, where the logarithm is taken on the mean number of solutions. The former is usually considered to be more representative, as unlikely samples leading to an abnormally large number of solutions can be observed to dominate the latter (recall Chap. 3), but requires a more involved calculation with the use of the replica trick. In the TAP case, quenched and annealed complexities coincide for solutions above a certain free energy threshold [106] (where most solutions lie but importantly not the ground state).

As a matter of fact, we have brought up that, even in the annealed case, the computation of the TAP complexity has proved to be a formidable task, and has sparked a large amount of controversy, as the original solution computed by Bray & Moore (BM) [106] has been put into question before being (partially) salvaged by the metastability of TAP states in the thermodynamic limit [109]. For a relatively up to date summary of the situation, we refer the reader to G. Parisi’s contribution in [214].

While the BM approach can be adapted to the NMFE [211,215], several aspects of the calculation remain unclear, particularly as the absence of a sub-dominant re-action term means that the argument of the metastability of states in the $N \rightarrow \infty$ limit is no longer valid *a priori*, although numerical results support the marginally




Figure 6.1: (a) Annealed complexity of the Naive Mean-Field Equation in (T, ε) space, where we recall $T = \beta^{-1}$, measured numerically for $N = 40$. The dashed line represents the critical temperature $T_c(\varepsilon)$ for which the paramagnetic fixed point ceases to be stable, while the continuous line indicates $\varepsilon = 1$ inferred from the $\beta \rightarrow \infty$ result. (b) Annealed complexity as a function of ε for varying temperatures $T < 1/\sqrt{2}$, i.e. in the bottom region of (a) where the complexity vanishes at $\varepsilon = 1$. The continuous line represents the $\beta \rightarrow \infty$ analytical solution, recovering the result of Tanaka and Edwards [112] $\Sigma \approx 0.1992$ for $\varepsilon = 0$. For $\varepsilon > 1$ and $T = 0$, the only possible fixed point is $m_i^* = 0, \forall i$.

stable nature of NMFE fixed points in the thermodynamic limit [211]. Its extension to $\varepsilon > 0$ is still ongoing.

Nevertheless, the previously introduced critical β_c and the existing computation of the number of fixed points as a function of ε in the $\beta \rightarrow \infty$ limit [201, 216] can be used to conjecture the boundaries of the region in (β, ε) space where the complexity Σ is non-vanishing. Indeed, in the zero temperature case, it has been shown [201] that the annealed complexity can be expressed as a function of the asymmetry parameter η defined in Eq. (5.10) as

$$\Sigma(\eta) = -\frac{1}{2}\eta x^2 + \log 2 + \log \Phi(\eta x), \quad (6.5)$$

with Φ the Gaussian cumulative density $\Phi(x) = \frac{1}{2} \operatorname{erfc}(-\frac{x}{\sqrt{2}})$ and x is the solution to

$$x\Phi(\eta x) = \Phi'(\eta x). \quad (6.6)$$

The main insight provided by this result is that the complexity vanishes at $\eta = 0$, corresponding to $\varepsilon = 1$, where the paramagnetic fixed point is supposed to be unstable as $\beta_c(\varepsilon = 1) = \sqrt{2}$. As the complexity is a decreasing function of temperature, this therefore means that $\varepsilon = 1$ is an upper limit for the existence of fixed points when β is finite. This conjecture is also consistent with the breakdown of fixed point solutions to the dynamical mean-field theory below the critical noise level that will be discussed in Sec. 6.5.1, as well as the saddle point equations obtained when adapting the BM calculation to $\varepsilon > 0$ (not presented here).

Combining these two somewhat heuristic delimitation for the existence of a large number of non-trivial fixed points, we obtain the critical lines shown in

Fig. 6.1(a). Overlaying these borders with the annealed complexity measured numerically following the method of [215], we find a very good agreement. In particular, the vanishing of the complexity at $\varepsilon = 1$ appears to be consistent for $T < T_c(\varepsilon = 1) = 1/\sqrt{2}$, as shown in Fig. 6.1(b). The agreement with the $\beta \rightarrow \infty$ analytical result, represented by the continuous line, also appears to validate our counting method at low temperatures. Note that one can in fact show that for $\varepsilon > 1$ and $N = \infty$, the only fixed point (or Nash equilibrium) is the “rock-paper-scissors” equilibrium $m_i^* = 0, \forall i$.

Let us finally briefly discuss the zero-temperature annealed complexity Σ when a cooperative contribution $J_0 > 0$ is present in Eq. (5.11), and in the special case $\varepsilon = 0$. As detailed in Appendix C.2, the difficulty with this computation is that the spin-glass solution $M = 0$ remains a valid fixed point for all values of J_0 . While we expect that it becomes unstable to single-flip perturbations for $J_0 > J_{0c}$, this stability criterion is not straightforward to implement analytically. Looking for solutions to the saddle point equations numerically, we find a solution associated to a mixed ($M > 0$) phase for $J_0 > J_{0c} \approx 1.3$. Importantly, the complexity in this mixed phase is markedly non-zero, although it decreases with J_0 as the solution approaches the ferromagnetic fixed point $M = 1$. In any case, we therefore expect most of the $J_0 = 0$ results to hold for $J_0 < J_{0c}$, but also perhaps in this mixed region where the complexity remains non-zero.

6.2 Counting limit cycles

In the previous section, we have established the region of parameter space where exponentially numerous fixed points exist, which might possibly be reached by learning in the slow limit $\alpha \ll 1$. However, limit cycles of various lengths turn out to also be exponentially numerous when $\varepsilon < 1$, so we need to discuss them as well before understanding the long term fate of the learning process within our stylized complex world.

6.2.1 Cycles without memory

In this memory-less limit, the dynamics becomes that of the extensively studied Hopfield model [35, 216, 217] where the binary variable represents the activation of a neuron evolving as

$$S_i(t+1) = \text{sign} \left(\sum_j J_{ij} S_j(t) \right), \quad (6.7)$$

with *parallel* updates. Counting limit cycles of length L is even more difficult than counting fixed points (which formally correspond to $L = 1$). Some progress

have been reported by Hwang *et al.* [201] in this memory-less case $\alpha = 1$. The notion of fixed point complexity Σ (defined in Eq. (6.4)) can be extended to limit cycle complexity Σ_L for limit cycles of length L , with $\Sigma_{L=1} \equiv \Sigma$. The results of Ref. [201] can be summarized as follows:

- When $\varepsilon < 1$, limit cycles with $L = 2$ have the largest complexity, which is exactly twice of the fixed point complexity: $\Sigma_2 = 2\Sigma_1$ (as was in fact previously shown by Gutfreund *et al.* [216]).
- The complexities $\Sigma_L(\varepsilon)$ all go to zero when $\varepsilon = 1$.
- When $1 < \varepsilon \leq 2$, limit cycles with $L = 4$ dominate, with $\Sigma_4(\varepsilon) \geq \Sigma_2(2 - \varepsilon)$.
- Close to $\varepsilon = 1$, the cut-off length L_c , beyond which limit cycles become exponentially rare, grow exponentially with N : $L_c \sim e^{aN}$, where a weakly depends on ε .

From this analysis, one may surmise that:

- When a limit cycle is reached by the dynamics, it is overwhelmingly likely to be of length $L = 2$ for $\varepsilon < 1$ and of length $L = 4$ for $1 < \varepsilon \leq 2$.
- Even if exponentially less numerous, exponentially long cycles will dominate when $e^{aN} > e^{N\Sigma_2}$, which occurs when $\varepsilon_c < \varepsilon < 2 - \varepsilon_c$, with $\varepsilon_c \approx 0.8$.

These predictions are well obeyed by our numerical data, recall Fig. 5.3(b) and see Fig. 6.3 below. Note however the strong finite N effects that show up in the latter figure, which we will be at the heart of the next section.

6.2.2 Cycles with memory

When $\alpha < 1$ and $\beta = \infty$, the update of $S_i(t)$ have the same fixed points independently of α , but of course different limit cycles, which may in fact cease to exist when α is small.

Here, we attempt to enumerate the number of cycles of length L , in the spirit of the calculation of Hwang *et al.* [201] for $\alpha < 1$. As detailed in Appendix C.3, we write the number of these cycles as a sum over all possible trajectories of a product of δ functions ensuring the $\alpha < 1$ dynamics of Q_i are satisfied between two consecutive time-steps, while a product of Heaviside step functions enforces $S_i(t) = \text{sign}(Q_i(t))$. Introducing the integral representation of the δ function, averaging over the disorder and taking appropriate changes of variable to decouple the N dimensions, the (annealed) complexity of cycles of length L writes

$$\Sigma_L(\alpha, \eta) = \text{saddle}_{\hat{R}, \hat{K}, \hat{V}} \left\{ \sum_{s < t} i\hat{R}(t, s)i\hat{K}(t, s) - \frac{\eta}{2} \sum_{t, s} \hat{V}(t, s)\hat{V}(s, t) + \log \mathcal{I}_L \right\}, \quad (6.8)$$




Figure 6.2: (a) $L = 2$ cycle complexity from the numerical resolution of the saddle point equations for $\varepsilon = \{0, 0.2, 0.4, 0.6\}$ from dark purple to light green, dashed lines indicating the fixed point complexity associated to each parameter. (b) Overlap between successive steps $C(\tau = 1) = -i\hat{R}(t, t + 1)$ at the $L = 2$ saddle, showing the nontrivial coalescence of the cycle and fixed point saddle point solutions as α is decreased. The numerical resolution appears to breakdown when we get close to $\alpha = 0$.

where $\hat{R}(t, s)$ and $\hat{K}(t, s)$ are symmetric matrices while $\hat{V}(t, s)$ is not *a priori*, and the expression of $\mathcal{I}_L$ is given in Appendix C.3. As a sanity check, one can verify that the $L = 1$ case, corresponding to the fixed point complexity, is indeed independent of α and is given by the same expression as Eq. (6.5), see Appendix C.3.1. In a similar vein, one can recover $\Sigma_2(\alpha = 1) = 2\Sigma(\eta)$ for all values of ε (Appendix C.3.2).

In the case of $L = 2$ cycles for $\alpha < 1$, the six coupled nonlinear saddle equations given in Appendix C.3.2 can be solved numerically. As shown in Fig. 6.2, it appears that $\Sigma_2(\alpha, \eta) \rightarrow \Sigma(\eta)$ when $\alpha \rightarrow 0^+$, while $C(t, t + 1) = -i\hat{R}(t, t + 1)$ can be observed to continuously tend to unity. While numerical difficulties prevent us from exploring very small values of α , it seems clear that the saddle point corresponding to the $L = 2$ cycles eventually coalesces with the fixed point saddle (which is known to be a sub-dominant saddle point when $\alpha = 1$, see [201]). In any case, and perhaps surprisingly, there does not appear to be a critical value of α below which fixed points become more abundant than cycles. We therefore expect a progressive crossover and not a sharp transition. As performed in [201] in the $\alpha = 1$ case, it would be interesting to study the complexity of $L = 2$ cycles for a given value of the shifted overlap $C(t, t + 1) = \frac{1}{N} \sum_i S_i(t) S_i(t + 1)$ in the spirit of a Franz-Parisi potential calculation [218] in order to rigorously establish the coalescence of the fixed point and oscillating saddles.

6.3 Dynamical Mean-Field Theory

In the previous section, we have seen that both fixed points and limit cycles are exponentially numerous. However, the question remains as to what happens *dynamically*, as the existence of a large number of fixed points or limit cycles by




Figure 6.3: Steady-state two-point correlation between configurations shifted by $\tau = 2$ time-steps in the $\alpha = 1$, $\beta \rightarrow \infty$ limit from finite N numerical simulations averaged over 128 samples of disorder and initial conditions, error bars showing 95% confidence intervals. The $N = \{64, 128\}$ simulations are run for $t = 10^8$ time-steps to illustrate taking the $t \rightarrow \infty$ limit before $N \rightarrow \infty$, whereas $N = \{256, 512, 1024\}$ have been simulated for $t = 5 \times 10^6$ time-steps to recover the $N \rightarrow \infty$ before $t \rightarrow \infty$ regime. The continuous line represents the $N \rightarrow \infty$ DMFT solution integrated numerically, while the vertical dotted line corresponds to the critical value ε_c found by Hwang *et al.* [201].

no means guarantees that these will be reached at long times.

In fact, the number of agents N is expected to play a major role in determining the long term fate of the system. In particular, there are strong indications that the time τ_r needed to reach a fixed point or a limit cycle grows itself exponentially with N , at least when $\alpha = 1$ [217]. More precisely,

$$\tau_r \sim N^s e^{NB(\varepsilon)}, \quad (6.9)$$

where s is an exponent (possibly dependent on ε) and $B(\varepsilon)$ an effective barrier such that $B(\varepsilon = 0) = 0$. Hence one expects that as N grows, fixed points/limit cycles will in fact never be reached (except for $\varepsilon = 0$), even if they are numerous. This is in fact what happens numerically, see Fig. 6.3.

In this large N regime, new types of behavior therefore appear, that one can call quasi-fixed points or quasi-cycles. In the case of quasi-fixed points, learning does not strictly speaking converge (it would take an infinite time to do so), but actions $S_i(t)$ fluctuate around fixed, time-independent averages. In other words, the two-point correlation $C(\tau)$ is not equal to one for all τ (which would be the case for a fixed point) but reaches a positive plateau value for large τ : $C(\tau \rightarrow \infty) = C_\infty > 0$. The same holds for quasi-cycles of length L if one considers the correlation function computed for $\tau = nL$, with n an integer: $C(nL \rightarrow \infty) = C_\infty > 0$. This being said, we will see that the schematic phase diagram drawn in Fig. 5.1 of the previous chapter continues to hold qualitatively in the large N finite t limit, provided one interprets “fixed points/cycles” as “quasi-fixed points/cycles” in the sense defined above.

In order to study the complicated learning dynamics that takes place in the true $N \rightarrow \infty$ limit, we will resort to Dynamical Mean-Field Theory (DMFT). In a nutshell, DMFT allows deterministic or stochastic dynamics in discrete or continuous time of a large number N of interacting degrees of freedom to be rewritten as a one-dimensional stochastic process with self-consistent conditions. While difficult to solve both analytically and numerically due to their self-consistent nature, DMFT equations have proved very effective at describing a very wide range of complex systems – see [219] for a recent review. Note however that such an approach is only valid when $N \rightarrow \infty$; as will be clear later, strong finite size effects can appear and change the conclusions obtained using DMFT.

6.3.1 Derivation

In our case, we write the DMFT for the evolution of the incentives $Q_i(t)$, which directly yield $m_i(t) = \tanh(\beta Q_i(t))$. In order to do so, we rewrite our *online* learning process, which depends on the realized $S_i(t)$, as an expression solely in terms of $m_i(t)$ with additional fluctuations,

$$\sum_j J_{ij} S_j(t) = \sum_j J_{ij} m_j(t) + \eta_i(t), \quad \eta_i(t) = \sum_j J_{ij} \xi_j(t), \quad (6.10)$$

with $\xi_i(t) = S_i(t) - m_i(t)$ and hence $\langle \xi_i(t) \rangle = 0$ and $\langle \xi_i(t) \xi_i(s) \rangle = (1 - (m_i(t))^2) \delta_{t,s}$. Now, assuming the central limit theorem holds, the random variables η_i become Gaussian for large N with

$$\langle \eta_i(t) \rangle = 0, \quad \langle \eta_i(t) \eta_j(s) \rangle = v(\varepsilon) (1 - q(t)) \delta_{t,s} \delta_{i,j}, \quad (6.11)$$

where $q(t) = C(t, t)$ as defined in Eq. (5.15). As required, in the noiseless limit $\beta \rightarrow \infty$ limit, one has $q(t) = 1 \forall t$ and the random variables η_i are identically zero.

Starting from the N equations

$$Q_i(t+1) = (1 - \alpha) Q_i(t) + \alpha \sum_j J_{ij} m_j(t) + \alpha \eta_i(t) + \alpha h_i(t), \quad (6.12)$$

where $h_i(t)$ is an arbitrary external field that will eventually be set to 0, the DMFT can be derived using path integral techniques or the cavity method, the latter being detailed in Appendix C.4. Remaining in discrete time to explore the entire range of values of α , one finds, in the $N \rightarrow \infty$ limit,

$$Q(t+1) = (1 - \alpha) Q(t) + \alpha^2 (1 - \varepsilon) \sum_{s < t} G(t, s) m(s) + \alpha \phi(t) + \alpha h(t), \quad (6.13)$$

with $\langle \phi(t) \rangle = 0$, and

$$\langle \phi(t) \phi(s) \rangle = v(\varepsilon) [C(t, s) + (1 - q(t)) \delta_{t,s}], \quad (6.14)$$

Chapter 6. Detailed analysis of the model

The memory kernel G and correlation function C are then to be determined self-consistently,

$$G(t, s) = \left\langle \frac{\partial m(t)}{\partial h(s)} \Big|_{h=0} \right\rangle, \quad C(t, s) = \langle m(t)m(s) \rangle, \quad (6.15)$$

where the averages $\langle \dots \rangle$ are over the realizations of the random variable ϕ .

While we shall see that this discrete description will allow for the numerical resolution of the $N \rightarrow \infty$ dynamics (see below), providing precious insights and intuitions, a continuous description will be much more convenient to obtain analytical insights. In the $\alpha \ll 1$, $t \gg 1$ regime, we can rescale the time as $t \rightarrow t/\alpha$. Interestingly, doing so requires expanding $Q(t+1)$ to the second order if one is to keep an explicit dependence on α . The resulting continuous dynamics reads

$$\frac{\alpha}{2} \ddot{Q}(t) + \dot{Q}(t) = -Q(t) + (1 - \varepsilon) \int_0^t ds G(t, s) m(s) + \phi(t) + h(t) \quad (6.16)$$

with

$$\langle \phi(t)\phi(s) \rangle = v(\varepsilon)[C(t, s) + \alpha(1 - q(t))\delta(t - s)], \quad (6.17)$$

and the memory kernel and correlation function are similarly defined self-consistently

$$G(t, s) = \left\langle \frac{\delta m(t)}{\delta h(s)} \Big|_{h=0} \right\rangle, \quad C(t, s) = \langle m(t)m(s) \rangle, \quad (6.18)$$

with, we recall, $q(t) = C(t, t) = \langle m^2(t) \rangle$. Very importantly, note the rescaling in time introduces a prefactor α in the variance of the ϕ , which stems from the noise in the learning process. Since $1 - q(t) \sim \beta^{-2}$ for large β , this extra term is of order α/β^2 , as anticipated above.

In the next sections, the DMFT equations will be used to shed light on the dynamical behavior of the model in the limit $N \rightarrow \infty$.

6.3.2 Numerical integration

As mentioned above, the discrete time dynamics can first be integrated numerically to obtain solutions free of any finite size effects [220]. Due to the self-consistent nature of the equations, their numerical resolution is not trivial, and requires an iterative scheme to update both the memory kernel and correlation function until convergence.

The procedure is essentially the same as that described in Ref. [221]. We start from an initial guess for both the correlation matrix and the memory kernel, we typically choose $C(t, s) = G(t, s) = \exp(-|t - s|)$, mostly for a lack of better idea. A realization of the noise trajectory is then drawn and Eq. (6.13) is iterated in time from a random initial condition $Q(0)$ and up to a fixed final time t_f . This procedure




Figure 6.4: Convergence of some selected observables with the iteration number i in the numerical resolution of the discrete DMFT equations constructed from $M = 5 \times 10^5$ independent trajectories for (a) $\alpha = 0.9$, $\beta \rightarrow \infty$, $\varepsilon = 0.4$; (b) $\alpha = 0.1$, $\beta = 4$, $\varepsilon = 0.6$. Note that the number of iterations required depends on the parameters.

is repeated independently a large number of times M (typically $M \sim 10^5$). Taking averages over this large number of realizations, the correlation matrix and memory kernel are updated following Eq. (6.15). Note that, as recommended in [221], we perform “soft” updates on $C(t, s)$ and $G(t, s)$, i.e. $X^{\text{updated}} = (1 - a)X^{\text{prev.}} + aX^{\text{new}}$. We found $a = 0.5$ to be appropriate in most cases. Examples of the convergence of different observables are shown in Fig. 6.4.

As detailed in the original work of Eissfeller & Oppen [222], we use Novikov’s theorem to compute the response function with correlations in order to avoid the unpleasant task of taking finite differences on noisy trajectories, at the cost of the inversion of the correlation matrix. This inversion will however mean that very long trajectories become difficult to integrate. In practice, we will often be limited to $t_f \leq 1000$, which we found to be sufficient in most cases.

6.3.3 Interpretation in a socioeconomic context

Before jumping into the study of the DMFT equations, let us take a sidestep to discuss the socioeconomic interpretation of Eq. (6.16) (or Eq. (6.13) for that matter). Remarkably, we have managed to reduce N interacting agents to a single equation for the incentive, or equivalently the intention – at the cost of self-consistency of course.

As hinted in a footnote of the introductory chapter, this is, in a sense, the equation describing a truly representative agent. Importantly, and at the heart of the difficulties we will encounter in studying this equation, these dynamics are non-Markovian, as they depend on the entire past trajectory through the memory kernel. The representative agent therefore *a priori* has long term memory, and aggregates past information in a much more complex way than any given agent in the system.

This is in fact the important conclusion that we can draw from this description of the model. While the representative agent paradigm is not inherently wrong, as one can construct an effective representative agent as we did, it is essential to notice that this agent does not correspond to *any* agent in the system: it is a new entity with markedly different features to those of the original agents. With this difference in mind, and as suggested in [223], DMFT might be a promising avenue to rejuvenate and redefine the representative agent paradigm in complex systems.

6.4 Noiseless learning

In this section, we use both the DMFT equations and the results on the complexity of fixed points and limit cycles to classify the different dynamical behaviors of the learning process in the noiseless case $\beta \rightarrow \infty$, where the realized and expected decisions are equal, $m_i(t) = S_i(t) = \text{sign}(Q_i(t))$.

6.4.1 The memory-less limit

In this case, corresponding to Eq. (6.7), both approaches (DMFT and complexity of limit cycles) seem to agree on the overall picture: as ε increases from 0 to 1, the system transitions from $L = 2$ cycles to over-damped oscillations and chaos – see Fig. 5.3 of the previous chapter. However, upon scrutiny, one realizes that the perfect agreement between DMFT and direct numerical simulations of the dynamics for finite N is only valid in a region where ε is small and N large enough – see Fig. 6.3. In particular, when $0.5 \lesssim \varepsilon \lesssim 0.8$, $L = 2$ cycles do persist when N is smaller than ~ 200 . For larger N , the lag 2 autocorrelation function $C(\tau = 2)$ is noticeably smaller than unity, and well predicted by DMFT as soon as $N \gtrsim 1000$.

What happens for $\varepsilon \lesssim 0.5$ when $N \rightarrow \infty$? The numerical solution of the DMFT equations suggest the following scenario: when $\varepsilon < \varepsilon_{\text{RM}} \approx 0.473$, the long time value m_∞ of the correlation with the initial conditions $C(0, 2n)$ at even time steps is strictly positive.³¹ In this regime, the value of $C(t, t + 2) \rightarrow C(\tau = 2)$ is only exactly equal to one for $\varepsilon = 0$ (permanent oscillations), and appears to reach a plateau close to unity for $\varepsilon > 0$. Below ε_{RM} , we can conclude that the system is not ergodic, which will have important implications on the finite temperature dynamics. For asymmetries greater than ε_{RM} on the other hand, the decorrelation becomes exponential and we enter a *bona fide* chaotic, ergodic regime [222].

Although memory-less learning is clearly unrealistic, these results are rather instructive. The system is indeed unable to display aggregate coordination when

³¹The convergence to m_∞ is as slow a power law of τ , which makes difficult its numerical determination. Obtaining the precise value of ε_{RM} is therefore challenging [222].

interactions are mutually independent (chaotic region around $\varepsilon = 1$). Placing ourselves in the socioeconomic setting, it seems evident that the number of players vastly exceeds the number of iterations, and the results above indicate that the chaotic region is in fact quite large. Perhaps more importantly, in the absence of learning, agents will see their decisions vary at a high frequency without ever reaching any static steady-state, not only when the game is close to zero sum ($\varepsilon \rightarrow 2$), but even when it is fully reciprocal ($\varepsilon \rightarrow 0$). Clearly, this last point underlines the importance of introducing memory to recover realistic learning dynamics.

6.4.2 Memory helps convergence to fixed points

For N not too large, we observe numerically that the fraction of “frozen” agents for which $S_i(t+1) = S_i(t)$ quickly tends to 1 as α decreases from 1, as shown in Fig. 6.5. This is somewhat consistent with intuition, as the learning dynamics average rewards over a period $\tau_\alpha \sim 1/\alpha$, meaning that high frequency cycles observed for $\alpha = 1$ are expected to be “washed out” when α is sufficiently small. Since fixed points exist in large numbers, it appears natural that they are eventually reached given their abundance at zero temperature. However, as we have shown in the previous section, $L = 2$ limit cycles are still much more numerous than fixed points for $\alpha \gtrsim 0.5$. The fact that $C(\tau = 1)$ approaches unity as α is reduced much faster than in Fig. 6.2(b) suggests that the basin of attraction of fixed points quickly expands, at the expense of $L = 2$ limit cycles.

Our numerical results therefore indicate that for any finite size system which has enough time to reach a steady state, the effect of α is effectively to help the system find fixed points – see Fig. 6.6. Focusing for example on the points corresponding to $N = 128$ and $\alpha = 0.1$, we indeed observe that the 95% quantile includes $C(2/\alpha) = 1$ even for $\varepsilon = 1$, i.e. fixed points can be reached even in the chaotic regime with modest simulation times, which would be an overwhelmingly improbable scenario in the memory-less case, as illustrated by Figs. 6.6 (b) and (c).

As the number of agents N increases, we enter the DMFT regime shown as plain lines in Fig. 6.6. One finds that decreasing the value of α slows down the decorrelation of the system. However, for small α , the evolution becomes a function of $\alpha\tau$ only, as suggested by Eq. (6.16) when $\alpha \rightarrow 0$: the dynamical slowdown is dominated by the long memory of learning itself.

Fig. 6.6 shows that when $\varepsilon \gtrsim 0.5$, sufficiently large systems (described by DMFT) decorrelate with time for all α , and we expect $C(\tau \rightarrow \infty) \rightarrow 0$ – learning also leads to chaos. When $\varepsilon \lesssim 0.5$, on the other hand, we found that there is ergodicity breaking, in the sense that $C(\tau \rightarrow \infty) > 0$, as we found above for cycles when $\alpha = 1$. More precisely, a numerical analysis of the DMFT equations




Figure 6.5: Convergence to fixed points with decreasing α for $\varepsilon < 0.8$, $\beta \rightarrow \infty$ and finite N . (a) Steady-state two-point correlation function between successive configurations from finite N numerical simulations averaged over 200 samples of disorder and initial conditions, error-bars showing 95% confidence intervals. (b), (c) and (d) Sample trajectories of 32 randomly chosen sites among $N = 256$ for $\varepsilon = 0.4$, $t_0 = 10^5/\alpha$, for $\alpha = \{1.0, 0.88, 0.7\}$ respectively.




Figure 6.6: Influence of finite size N , non-reciprocity ε and memory span α on the learning dynamics. (a) Steady-state two-point correlation function shifted by $\tau = 2/\alpha$ in the $\beta \rightarrow \infty$ limit for different memory loss rates α , from light green to black (color map on the right axis). Symbols correspond to direct simulations and plain lines to the solutions of the DMFT equations. The $N = 128$ simulations are initialized with $t_0 = 10^8$ time-steps, whereas $N = \{256, 512\}$ have been simulated for $t_0 = 10^6$ iterations before taking measurements. Results are averaged over 32 samples, with error-bars showing 95% confidence intervals. (b) and (c) Sample trajectories for all $N = 128$ sites for $\varepsilon = 1$ and $\alpha = \{1.0, 0.1\}$ respectively, clearly reaching chaotic and fixed-point steady-states.

suggests that when $\varepsilon, \alpha \rightarrow 0$, $C(\tau \rightarrow \infty) = 1 - c\sqrt{\varepsilon}$. We have not been able to find an analytical proof for such a singular dependence on ε .

In other words, there again seems to exist a critical value ε_{RM} separating the ergodic, chaotic phase for $\varepsilon > \varepsilon_{\text{RM}}$ from the non-ergodic, quasi fixed-point behavior for $\varepsilon < \varepsilon_{\text{RM}}$. However, our numerical results are not precise enough to ascertain the dependence of ε_{RM} on α , which seems to hover around the value 0.473 found for $\alpha = 1$. More work on this specific point would be needed to understand such a weak dependence on the memory length.

The precise dynamical behavior of the autocorrelation function $C(\tau)$ can be ascertained in the continuous limit $\alpha \rightarrow 0$ when $\varepsilon = 1$. Indeed, the influence of the memory kernel vanishes in this case where interactions are exactly non-symmetric, leaving us with

$$\dot{Q}(t) = -Q(t) + \phi(t), \quad (\alpha \rightarrow 0) \quad (6.19)$$

where we emphasize that the time variable has been rescaled as $t \rightarrow \alpha t$. From there, the classical solution method proposed by Sompolinsky & Crisanti [224] can be adapted with a small modification due to our parametrization of the interaction matrix that scales the variance of the entries by a factor $1/2$ for $\varepsilon = 1$. The result is an exponentially decreasing autocorrelation $C(\tau) \approx e^{-\tau/\tau_1}$, with relaxation time

$$\tau_1 = \sqrt{\frac{2}{1 - 2/\pi}}. \quad (6.20)$$

This theoretical result is compared with simulations and the numerically integrated DMFT equations for specific values of α in Fig. 5.3, displaying a very satisfactory match, even for relatively large values of α . It should be noted that the memory in the system seems to induce strong finite-size effects. As visible in the right figure, the long time autocorrelation remains clearly above the theoretical value for N as high as 1024, and appears to only slowly vanish when α gets larger.

6.4.3 Anomalous stretching of cycles

We have seen that understanding how memory allows the system to find fixed-points, or quasi fixed-points, when they exist is a challenging task. While the very precise behavior of the correlation function is difficult to ascertain, we have nonetheless obtained a reasonable picture of the role of the memory loss parameter α . A central question is now what will happen if the memory loss rate is reduced in the region of parameter space where there are only limit cycles. As previously stated, averaging over a period $\tau_\alpha \sim 1/\alpha$ clearly suggests that the occurrence of short cycles (starting at $L = 4$ for $\alpha = 1$) should gradually vanish.




Figure 6.7: Evolution of the time shifted autocorrelation function in the non-symmetric case $\varepsilon = 1$, $\beta \rightarrow \infty$, for different system sizes and memory loss parameters averaged over 20 realizations. (a) lin-lin scale, errorbars showing 95% confidence intervals, continuous lines representing the numerically integrated DMFT equations; (b) lin-log scale, highlighting the exponential decay characteristic of chaos, and rescaling of the time shift by α such that points collapse onto the analytical solution of Sompolinsky *et al.* (continuous black line), errorbars not shown in the interest of readability. Note the visible finite size effects when the autocorrelation is close to 0.




Figure 6.8: Evolution of the oscillation frequency ω_0 with the memory loss rate α for $\beta \rightarrow \infty$, $\varepsilon > 2 - \varepsilon_c$ averaged over 96 samples of disorder and initial conditions, errorbars showing 95% confidence interval. (a) Log-log plot of ω_0 as a function of α , dashed line corresponding to $\omega_0 \sim \sqrt{\alpha}$. (b) Two-point autocorrelation for $\varepsilon = 1.8$, $N = 256$ as a function of the rescaled time lag for different values of α . (c) Power spectrum of the autocorrelation for the same parameters, displaying secondary peaks at odd multiples of this fundamental frequency, as expected from the triangular aspect of $C(\tau)$.

Naively, one might expect a simple rescaling in time $t \rightarrow t/\alpha$, yielding cycles – when they exist – of period inversely proportional to α itself. Looking at the numerical results from both the finite size game and the DMFT integrated numerically in Fig. 6.8(a), it quickly appears that such a simple rescaling in time does not provide the correct description. Indeed, the period of cycles is observed to be proportional to $1/\sqrt{\alpha}$, i.e. much shorter than $1/\alpha$ – see Figs. 6.8(b) and (c).

One important aspect to note is that there is some decorrelation, as the second peak of $C(\tau)$ does not quite reach unity (in Fig. 6.8(b)), meaning that we may see quasi-cycles and not exact limit cycles, complicating the analytical description of the phenomenon. Just as true fixed points in the $N \rightarrow \infty$ limit only exist only $\varepsilon = 0$, it appears that only the case $\varepsilon = 2$ does display true limit cycles.

Another subtle point to consider is similar to the $\varepsilon < 1$ cases discussed, we expect the time taken to reach these cycles will depend on the system size and the relative distance to the chaotic region. This is confirmed by the DMFT solved for fixed trajectory times for $\varepsilon = 1.5$ (light crosses), which progressively departs from the $\omega_0 \sim \sqrt{\alpha}$ regime around $\alpha = 0.1$.

To understand how such non-trivial stretching occurs, we go back to the continuous DMFT equation,

$$\frac{\alpha}{2}\ddot{Q}(t) = -\dot{Q}(t) - Q(t) + (1 - \varepsilon) \int_0^t ds G(t, s)m(s) + \phi(t) + h(t).$$

While the presence of the second order derivative $\ddot{Q}(t)$ appears natural to recover limit cycles, it should be noted that this term, being pre-factored by α , is superficially subdominant relative to the dissipation represented by $\dot{Q}(t)$. While we have seen that there is some decorrelation, the fact that robust oscillations are present therefore suggests that the complicated self-consistent forcing terms almost exactly compensate dissipation over a period, allowing the system to periodically revisit quasi-identical configurations. In fact, the shape of these oscillations is far from sinusoidal, but rather of see-saw type, see Fig. 6.8(b). This suggests that in the limit $\alpha \rightarrow 0$, $\ddot{Q}(t)$ diverges each time $\dot{C}(\tau)$ changes sign, such that $\frac{\alpha}{2}\ddot{Q}(t)$ cannot be neglected and therefore sets the relevant time scale to $\alpha^{-1/2}$. We have however not been able to perform a more precise singular perturbation analysis of this phenomenon.

Going beyond this rather loose argument, and precisely characterizing such see-saw patterns appears very challenging and is left for future work. A possible approach would be to first take the $\varepsilon = 2$ case where true cycles should exist, and to assume the correlation function is an exact triangular wave of frequency ω as suggested by Fig. 6.8(c). As a result, $m(s) = \text{sign}(Q(s))$ is an exact square wave, and the convolution with G can be written as a product in Fourier space. Enforcing the dissipation over a period to be zero, one could then perhaps find




Figure 6.9: Heat map of $q = \lim_{t \rightarrow \infty} C(t, t) = C(\tau = 0)$ in $(T = 1/\beta, \varepsilon)$ space from numerical simulations for $N = 256$, $t_0 = 10^6$ averaged over 32 samples of disorder and initial conditions and (a), (b) and (c) corresponding to $\alpha = \{0.5, 0.1, 0.01\}$ respectively. The white dashed line represents the critical temperature $T_c(\varepsilon)$ where the paramagnetic solution ($q = 0$) becomes linearly unstable (Sec. 6.1).

a closed equation for Q and ω if appropriate *ansatz* for the response and forcing functions are taken.

6.5 Noisy learning

While we have shown that the $\beta \rightarrow \infty$ deterministic limit can be relatively well understood with the tools at our disposal, one of the key features of our model is the uncertainty in the decision occurring for boundedly rational agents. Besides, it is also in this situation that the *online* learning dynamics differ significantly from the more widely studied *offline* learning where the entire model can be understood in terms of deterministic mixed strategies parameterized by the coefficients $m_i(t)$ (compare Eqs. (5.5) and (5.7) in the introduction of the model).

When α is close to unity and β becomes small, the fluctuations are too large for coordination to occur. Taking for instance $\alpha = 1$, it is indeed clear that the iteration

$$m_i(t+1) = \tanh \left(\beta \sum_j J_{ij} S_j(t) \right)$$

will have extremely large fluctuation in the argument on the right hand side. As a result, we expect to lose the sharp transition as a function of T that can be observed for the NMFE (see Fig. 6.9). The order parameter q instead continuously tends to 0 with T , regardless of the asymmetry ε . This regime is shown in Fig. 6.9(a), representing the heat map of $q = \lim_{t \rightarrow \infty} C(t, t)$ for $\alpha = 0.5$. Clearly, the linear stability analysis of the paramagnetic fixed point presented in Sec. 6.1 cannot hold when the thermal fluctuations are not averaged on large periods of time. To find a richer phenomenology, we will therefore focus on the $\alpha \ll 1$ regime where more complex dynamics can be observed.

6.5.1 (Quasi-)fixed points

In Sec. 6.1, we studied the fixed points of the NMFE that the game should reach if the fluctuations from imperfect learning can be neglected, i.e. if $\alpha \rightarrow 0$. Now, the DMFT equations that proved effective in the zero “temperature” limit can be used to revisit these finite β *quasi* fixed points. Indeed, going back to Eq. (6.16) and neglecting the term in $\sqrt{\alpha}$ from the correlation function as we did in the static setup, quasi-fixed points should satisfy

$$Q = (1 - \varepsilon)m\chi + J\sqrt{qv(\varepsilon)}z, \quad (6.21)$$

where z is now a *static* white noise of unit variance, q is simply the now constant autocorrelation and χ is the integrated response function that we assume to be time-translation invariant,

$$\chi = \int_0^\infty d\tau G(\tau). \quad (6.22)$$

The averages on the effective process can now be taken on z to self-consistently solve for q and χ (see e.g. [196] for a more detailed description). The resulting set of equations are then

$$q = \langle m^2(z) \rangle_z, \quad (6.23)$$

to be solved simultaneously with

$$\chi = \left\langle \frac{\beta(1 - m^2(z))}{1 - \beta(1 - \varepsilon)\chi(1 - m^2(z))} \right\rangle_z \quad (6.24)$$

where $m(z)$ is the solution to

$$m(z) = \tanh(\beta(1 - \varepsilon)\chi m(z) + \beta\sqrt{qv(\varepsilon)}z). \quad (6.25)$$

Although our model is entirely built on a dynamical evolution equation, and not on a notion of thermal equilibrium, this set of self-consistent equations coincides with the replica-symmetric solution of the NMFE model found by Bray, Sompolinsky & Yu [104] for $\varepsilon = 0$. Since replica symmetry is broken in the whole low temperature phase of the NMFE model, we expect that these static solutions of the DMFT cannot describe correctly the long time limit of the dynamics, as we now show.

The numerical solutions for the DMFT fixed point equations are shown in Fig. 6.10(a) and compared to numerical results of the game for small α and for $\varepsilon = 0.1$ and $\varepsilon = 0.6$. We find that the long time behavior of q for the direct simulation of the SK game (circles and squares) and for long time dynamical solution of the DMFT equations match very well, but differ from the value of q inferred from the set of self-consistent equations established above. This is expected since with such solution the order parameter q approaches unity exponentially fast as $T \rightarrow 0$,

Chapter 6. Detailed analysis of the model

whereas the fact that the probability of small local fields (i.e. rewards in the game analogy) vanish linearly (recall Fig. 5.5) suggest that $q = 1 - \kappa T^2$, as for the full RSB solution of [104] but with presumably a different value of κ . Remarkably, in this regime, the results from the long time dynamical solution to the DMFT equations (and the simulations) seem to collapse when plotted vs $T/T_c(\varepsilon)$, a result we have not been able to uncover analytically.

To ascertain the range over which this non-trivial mean-field solution should be valid, we can study the stability of the DMFT fixed point close to the critical temperature $1/\beta_c$, following the procedure first detailed in [195]. Considering a random perturbation to the fixed point $\epsilon \xi(t)$, with $\xi(t)$ a Gaussian white noise and $\epsilon \ll 1$, we study the perturbed solution

$$Q(t) = Q_0 + \epsilon Q_1(t), \quad (6.26)$$

with Q_0 the fixed point given in Eq.(6.21), where the noise is no longer static but similarly given by $\phi(t) = \sqrt{qv(\varepsilon)}z + \epsilon \phi_1(t)$. Replacing in the DMFT continuous dynamics for $\alpha \rightarrow 0^+$ and collecting terms of order ϵ , we find that the perturbation evolves as

$$\dot{Q}_1(t) = -Q_1(t) + \beta(1 - \varepsilon)(1 - m^2(z)) \int_0^t G(t, s) Q_1(s) + \phi_1(t) + \xi(t), \quad (6.27)$$

where we have used $\text{sech}^2(\beta Q_0) = 1 - m^2(z)$ from Eq. (6.25), giving in Fourier space

$$\hat{Q}_1(\omega) = \frac{\hat{\phi}_1(\omega) + \hat{\xi}(\omega)}{i\omega + 1 - \beta(1 - \varepsilon)(1 - m^2(z))\hat{G}(\omega)}, \quad (6.28)$$

where we have again assumed that the memory kernel is time-translation invariant.

Now, in the limit $\beta Q_1 \ll 1$, i.e. close to the critical temperature, one can linearize the hyperbolic tangent $\tanh(\beta Q_1(t))$ and write a closed equation for the spectral density of Q_1 at order $\beta^2 Q_1^2$,

$$\frac{1}{\langle |\hat{Q}_1(\omega)|^2 \rangle} = |i\omega + 1 - \beta(1 - \varepsilon)\hat{G}(\omega)|^2 - \beta^2 v(\varepsilon). \quad (6.29)$$

As a result, we have criterion for the onset of instability for $\omega = 0$:

$$(1 - \varepsilon)\chi = 1 - \beta\sqrt{v(\varepsilon)}, \quad (6.30)$$

where we noticed $\hat{G}(\omega = 0) = \chi$, given, close to β_c , by

$$\chi = \frac{1}{2\beta(1 - \varepsilon)} \left(1 - \sqrt{1 - 4\beta^2(1 - \varepsilon)} \right). \quad (6.31)$$




Figure 6.10: (a) Order parameter $q = C(t, t)$ averaged in time in the (quasi) stationary regime vs. rescaled temperature for $\varepsilon = \{0.1, 0.4\}$ (dark and light shades respectively). Circular and diamond markers correspond direct, finite size simulations at $N = 256$ for $\alpha = 0.01$ and $\alpha = 0.1$ respectively, whereas crosses represent the (dynamical) numerical solution to the complete set of $N \rightarrow \infty$ DMFT equations for $\alpha = 0.1$. Continuous lines show the solutions to the static DMFT fixed point equations (Eq. (6.23)-(6.25)). (b) Spectral density of a small perturbation Q_1 to the fixed point solution of the DMFT close to the critical temperature. As the quantity is necessarily positive for a valid solution, the grey region corresponds to instability.

Taking $\varepsilon = 0$, we recover the criterion found by Bray *et al.* [104] for the critical temperature, giving $T_c = 1/\beta_c = 2$ in their case.

For non-zero ε , we can also find the critical temperature by replacing Eq. (6.31) in (6.30), to recover yet again the critical temperature given by Eq. (6.3). As shown in Fig. 6.10(b), the spectral density evaluated at $\omega = 0$ becomes negative for $T < T_c(\varepsilon)$ for the fixed point solution. While the linearization used to obtain the relation is expected to become invalid, we understand this negativity as a strong sign that the solution is unstable and thus likely invalid, consistently with the discrepancy observed between the static prediction and the direct simulations and the dynamic solution of the DMFT show in Fig. 6.10(a).

6.5.2 Memory onset transition

The breakdown of the DMFT fixed point solution that is an average over not only realizations of the noise but also implicitly on initial conditions at the critical temperature is consistent with the emergence of a non-zero complexity for all $\varepsilon < 1$ at that point described in Sec. 6.1. Indeed, when there are an exponential number of fixed points to reach, it appears natural for the dynamics to no longer be time-translation invariant, contrary to what was postulated in the integrated response.

In the infinite time limit, the breakdown of the time translation invariant (TTI) nature of the response function can be studied following the procedure suggested

Chapter 6. Detailed analysis of the model

in [225]. Introducing a small TTI-breaking component to the memory kernel,

$$G(t, s) = G_0(t - s) + \epsilon G_1(t, s), \quad (6.32)$$

we can indeed identify the so-called memory onset transition, corresponding to the point in parameter space where a solution with $G_1(t, s) \neq 0$ becomes valid. Writing the fixed point equation with this memory kernel *ansatz* and self-consistently taking $G(t, u) = \left\langle \frac{\delta m(t)}{\delta h(u)} \Big|_{h=0} \right\rangle$, collecting terms of order ϵ and taking the limit $t \rightarrow \infty$, we have

$$G_1(u) = \beta(1 - \epsilon) \left\langle \frac{1 - m^2(z)}{1 - \beta^2(1 - \epsilon)\tilde{\chi}(1 - m^2(z))} \right\rangle_z \int_0^\infty ds G_1(s) G_0(s - u), \quad (6.33)$$

where $G_1(s) = \lim_{t \rightarrow \infty} G_1(t, s)$ is assumed to converge to a fixed value. Going once again to Fourier space and focusing on $\omega = 0$, we find that the TTI-breaking contribution may be non-zero only if we have

$$\beta^2(1 - \epsilon)\tilde{\chi}^2 = 1. \quad (6.34)$$

We notice that in the Hamiltonian case $\epsilon = 0$, the above *a priori* differs from the criterion defining the de Almeida-Thouless line for the NMFE [104]. Nonetheless, both are satisfied at $T = T_c$, the NMFE just as the standard SK model displaying replica symmetry breaking in the entirety of the spin-glass phase. In fact, above the critical temperature where the DMFT fixed points are valid and stable for $m(z) = 0$, one can check that both conditions are exactly equivalent.

When the interactions are non-reciprocal, however, the condition is no longer satisfied for $T \geq T_c$. Given the apparent invalidity of the fixed point solution at lower temperatures, we cannot formally determine if the memory onset eventually takes place at a smaller temperature for $\epsilon > 0$. Note, however, that even the complete absence of a memory onset transition does not mean we do not expect to observe a (partial) breakdown of ergodicity. Indeed, the memory onset transition is found by considering the violation of TTI at *infinite* times. As will become apparent in the next section, the long-time dynamics for $\epsilon > 0$ will still display aging effects for a wide range of temperatures.

6.5.3 Aging

The inadequacy of the static solution of the DMFT equations to describe the long term dynamics of the system is a well known symptom associated with the “aging” phenomenon [119], i.e. the fact that equilibrium is never reached and all correlation functions depend on the “age” of the system, recall Sec. 2.2.5:

$$C(t_w, t_w + t) = C_{\text{relax}}(t) + C_{\text{aging}}(t, t_w), \quad (6.35)$$




Figure 6.11: Aging behavior of the system for $N = 256$, averages performed over 576 realizations. Color indicates the value of αt_w , see scale on the far right. (a) and (b): Aging two-point correlation functions with the initial power law decay removed to isolate the aging component plotted as a function of t/t_w , inset showing the entire correlation function as a function of αt , dashed line representing the power law fit of the first relaxation. (a) $\alpha = 0.5$, $\beta = 4$, $\varepsilon = 0$, (b) $\alpha = 0.2$, $\beta = 2$, $\varepsilon = 0.25$. (c) and (d): Partially aging two-point correlation functions plotted as a function of αt . (c) $\alpha = 0.5$, $\beta = 4$, $\varepsilon = 0.5$, (d) $\alpha = 0.2$, $\beta = 2$, $\varepsilon = 0.5$.

where t_w is the waiting time, or age of the system, $t_w = 0$ corresponding to a random initial condition.

Aging typically arises in complex systems in the low temperature (low noise) limit. As described in Chap. 2, the energy landscape of such systems (like spin-glasses) are highly non convex and “rugged”, with a very large number of local minima or quasi-stable saddle points in which the dynamics gets stuck for extended periods of time [121]. Such *aging* phenomena are known to occur in a wide range of complex systems with reciprocal interactions, such as glassy systems, populations dynamics [226] or neural networks that are described by very similar mean-field dynamics [227]. Aging dynamics was also recently found in one of the aforementioned “habit formation” model, see [175].

Not surprisingly in view of its similarity with usual spin-glasses, the SK-game displays aging for reciprocal interactions ($\varepsilon = 0$) and sufficiently low temperatures $\beta > \beta_c$, see Fig. 6.11(a), for which Eq. (6.35) accurately describes the data with the initial relaxation component $C_{\text{relax}}(t)$ well fitted by a power law t^{-x} and

$$C_{\text{aging}}(t, t_w) = \mathcal{C} \left(\frac{t}{t_w} \right), \quad (6.36)$$

corresponding to the aging behavior found in a wide range of glassy models [97, 118–121, 228, 229].

What happens in our model when $\varepsilon > 0$? It is known from previous work (in somewhat different contexts), that aging survives in the presence of non-reciprocal

interactions, provided the asymmetry is not too large [230, 231]. If the asymmetry strength is further increased, we expect the amount of mixing in the system to eventually be large enough for the dynamics to no longer get stuck [232]. From our numerical simulations, shown in Fig. 6.11(b)-(d), one can conjecture that Eq. (6.35) still holds when $\varepsilon < \varepsilon^*$, but that the dynamics becomes time translation invariant when $\varepsilon > \varepsilon^*$. It is tempting to conjecture that aging disappears exactly when the dynamics becomes ergodic, i.e. when the correlation with a random initial condition decays to zero. This suggests that $\varepsilon^* = \varepsilon_{\text{RM}} \approx 0.47$, which is compatible with our numerical data. Note however that the transition between aging dynamics and time translation invariant correlations seems to occur somewhat progressively: the time required to reach a time-translation invariant regime is a decreasing function of not only ε but also the temperature $1/\beta$.

A particularity of the aging dynamics of the SK game is related to its *on-line* learning dynamics. As clearly visible in both the DMFT equation and the derivation of the Naive Mean-Field equation detailed in Appendix C.1, there will inevitably be some decorrelation in time of the *expected* decisions if $\alpha(1 - q)$ becomes significant. We therefore naturally expect the region where time translation invariance breaks down to be dependent on all three parameters α , β and ε .

The interpretation of aging in the socio-economic context is quite interesting and has been mentioned in Sec. 5.3.3. In a nutshell, it means that as time goes on, agents get stuck in locally satisficing strategies for longer and longer, but after a time proportional to the total time the game has already been played, the system eventually evolves and individual strategies m_i reach an altogether different configuration. This process goes on forever, but becomes slower and slower with time: the notion of quasi-equilibrium therefore makes sense at long times, for small enough noise and small enough non-reciprocity.

6.5.4 Chaos and (quasi-)limit cycles

When the non-reciprocity of interactions is sufficiently small and quasi-fixed points exist, we have established that boundedly rational systems display complicated aging dynamics when learning noise, parameterised by the value of β , is present. The immediate question is now how such noise influences the complex dynamics, chaos and (quasi) limit cycles that we have found in the $\beta \rightarrow \infty$, $\alpha \ll 1$ regime (see Sec. 6.4).

To qualitatively illustrate the effect of non-zero noise, we have run simulations for $\alpha = 0.1$ and different values of ε and β . Fig. 6.12 displays individual trajectories of the intentions $m_i(t)$, as well as the distribution of the values of individual m_i over all agents, realizations and time-steps. We also show the auto-correlation function $C(\tau)$ (assumed to be time-translation invariant over the short time scales considered), which is also compared to the DMFT solved numerically.




Figure 6.12: Finite β trajectories obtained from numerical simulations at $\alpha = 0.1$, $N = 256$, $t_0 = 10^8$ averaged over 96 samples of disorder and initial conditions for (a) $\beta = 4$, (b) $\beta = 2$ and from top to bottom $\varepsilon = \{0.1, 0.85, 1.05, 1.5\}$. Each line displays (from left to right) the evolution of 16 randomly selected agents, the associated histogram of m_i over both agents, time and realizations and the autocorrelation function assumed to be time translation invariant on short time scales.

Note that for the smallest value $\varepsilon = 0.1$, the trajectories illustrate the previously discussed quasi-fixed points emerging from the online dynamics. Clearly, while the correlation function remains close to constant, individual intentions are not exactly frozen (notice the wiggles in the top row of Fig. 6.12, specially for $\beta = 2$), explaining how the system as a whole eventually decorrelates and displays aging, as discussed in the previous section.

In the chaotic regime around $\varepsilon = 1$, the temperature has an immediate and expected effect on the order parameter. Decreasing β (increasing noise) obviously spreads the distribution of the m_i , which is less and less concentrated around ± 1 , as an immediate consequence of the smoothed out hyperbolic tangent. As a result, the equal-time autocorrelation $C(0) = q$ naturally decreases when β decreases. It is furthermore interesting to note that its value significantly decreases as the asymmetry parameter ε increases. We expect the role of α to be similar, as suggested by the phase diagrams presented in Fig. 6.9(b) and (c).

Dynamically, the decay of the autocorrelation appears to be independent of the strength of the noise.³² While it is known that external noise kills deterministic chaos in neural networks with uncorrelated couplings [233], what is interesting in our case is that both the non-linearity of the hyperbolic tangent (governed by β), and the strength of the effective noise (which scales as $1 - q$, see Eq. (6.17)) are varied simultaneously and, in a sense, self-consistently. Determining the way the decorrelation rate evolves with both β and ε is therefore quite non trivial and would be a very interesting endeavor.

Where we previously had limit cycles, for $\varepsilon = 1.5$ for instance (last row of Fig. 6.12), it appears that oscillations survive for large values of β . Note however that in this case the value $q = C(0)$ appears to very quickly vanish when β decreases. This is consistent with the linear stability analysis of the paramagnetic solution that should be valid for vanishingly small α . As visible in Fig. 6.9(c) we indeed expect that the region in which the system displays any form of aggregate coordination becomes increasingly narrow as ε gets closer to its maximum value of 2. Precisely for $\varepsilon = 2$, the system likely becomes fully disordered ($q = 0$) for any finite values of β when $\alpha \rightarrow 0$. For small but finite values of α as those presented here, this is not quite the case however, and large asymmetries $\varepsilon > 2 - \varepsilon_c$ do give rise to clear oscillations, both in individual trajectories and the correlation function (for instance Fig. 6.12, bottom row shows for $\varepsilon = 1.5$ that some oscillations can be somewhat sustained).

³²This is visible with a vertical log scale, for which the slope of the correlation function can be seen to be identical for different temperatures (not shown).

6.6 Conclusion

6.6.1 Blindsided by complexity

Let us once again summarize our main conceptual assertions, reiterating the messages of Chap. 5. As a schematic model of the complexity economic agents are confronted with, we introduced the “SK-game”, a discrete time binary choice model with N interacting agents and three parameters: α (memory loss rate), β (amount of noise in the learning process) and ε (non reciprocity of interactions).

We have shown that even in a completely static environment where the pay-off matrix does not evolve, agents are unable to learn collectively optimal strategies. This is either because the learning process gets trapped by a sub-optimal fixed point (or remains around one for very long times), or because learning never converges and leads to a never ending (chaotic or quasi-periodic) evolution of agents intentions.

Hence, contrarily to the hope that learning might save the “rational expectation” framework [172], which still holds the upper hand in macroeconomics textbooks, we argue that complex situations are generically *unlearnable*. Agents, therefore, must do with *satisficing* solutions, as argued long ago by H. Simon [63], an idea embodied by our model in a concrete and tangible way.

Only a centralized, omniscient agent may be able to ascribe an optimal strategy to all agents – which incidentally raises the question of trust: would agents even agree to follow the central planner advice? Would they even believe in her ability to solve such a complex problem, knowing that the solution sensitively depends on all parameters of the model? If a finite fraction of all agents fail to comply, the resulting average reward will drop precipitously below the optimal value and not be much better than the result obtained through individual learning.

As general ideas of interest in a socio-economic context, we have established that

1. long memory of past rewards is beneficial to learning whereas over-reaction to recent past is detrimental;
2. increased competition generically destabilizes fixed points and leads first to chaos and, in the high competition limit, to quasi-cycles;
3. some amount of noise in the learning process, quite paradoxically, allows the system to reach better collective decisions, in the sense that the average reward is increased;
4. non-ergodic behavior spontaneously appear (in the form of “aging”) in a large swath of parameter space, when α , $1/\beta$ and ε are small.

On the positive side, we have shown that learning is far from useless: instead of getting stuck among one of the most numerous fixed points with low average reward, the learning process does allow the system to coordinate around *satisficing* solutions with rather high (but not optimal) average reward. Numerically, the average reward at the end of the learning process is, for $\varepsilon = 0$, $\approx 8\%$ below the optimal value, when the majority of fixed points lead to an average reward equal to $2/3$ of the optimal value [106].

6.6.2 Technical results and conjectures

From a statistical mechanics perspective, our model is next of kin to several well studied models; a synthesis of our original results can be found in Figs. 5.1, 5.2, 5.4, 6.1, 6.9.

For example, when $\alpha = 1$, $\beta \rightarrow \infty$ limit, the dynamics is equivalent to a Hopfield model of learning with non-symmetric Gaussian synaptic couplings, for which many results are known, in particular on the number of fixed points and L -cycles. Introducing some memory with $\alpha < 1$, we found that the short limit cycles in which the $\alpha = 1$ parallel dynamics get stuck appear to progressively merge with fixed points. We also showed how the number of L -cycles can in principle be calculated for all values of $\alpha < 1$.

The chaotic region that is known to exist when interactions are mostly non-symmetric ($\varepsilon \approx 1$) also appears to be reduced by memory. When couplings are mostly non-reciprocal $\varepsilon \lesssim 2$, periodic oscillations survive but we found that decreasing α non-trivially increases the cycle length, as $\alpha^{-1/2}$.

When β is finite, the fixed point solutions to the dynamics correspond to the so-called Naive Mean-field Equation of spin glasses, another model that has been studied in detail [104]. One knows in particular that such solutions become exponentially abundant for small enough noise $1/\beta$ and for $\varepsilon = 0$, a result that we have extended to all $\varepsilon < 1$. When $\varepsilon > 1$, on the other hand, the only fixed point (or Nash equilibrium) is $m_i = 0$, $\forall i$.

The Dynamical Mean-Field Theory (DMFT) is a tool of choice for investigating the dynamics of the model when $N \rightarrow \infty$. DMFT is however frustratingly difficult to exploit analytically in the general case, so we are left with numerical solutions of our DMFT equations that accurately match direct numerical simulations of the model when N is large, but fails to capture some specific features arising when N is small. From our numerical results, we conjectured that quasi-fixed points (and correspondingly, aging dynamics) persist for small noise and when $\varepsilon \leq \varepsilon^*$, where the value of ε^* is difficult to ascertain but could be as high as $\varepsilon_{\text{RM}} = 0.47$, perhaps related to the remnant magnetisation transition found in [222].

For $\beta = \infty$, the long-time, zero temperature autocorrelation $C(\infty)$ appears to drop anomalously from the (assumed) value 1 at $\varepsilon = 0$, but we have no analytical

proof. DMFT equations also clearly lead to the anomalous $\alpha^{-1/2}$ stretching of the cycles mentioned above, but an analytic solution again eluded us.

One of the reason analytical progress with DMFT is difficult is presumably related to the phenomenon of “Replica Symmetry Breaking” discussed in Sec. 2.2 and its avatar in the present context of dynamical learning. Indeed, any attempt to expand around a static solution of the DMFT equations leads to inconsistencies in the interesting situation $\beta > \beta_c(\varepsilon)$ when decisions are not purely random, see Fig. 6.10(b). In fact, as shown in Fig. 6.10(a), the value of the order parameter q predicted by such static solutions is substantially off the value found from the long time, numerical solution of the DMFT equations, which itself coincide with direct simulations of the SK-game. *En passant*, we noticed that the value of q seems to be given by a universal function of $\beta_c(\varepsilon)/\beta$, independently of the value of ε . Again, we have not been able to understand why this would be the case.

Finally, we have numerically established several interesting results concerning average and individual rewards (Chap. 5), that would deserve further investigations. For example, the average reward seems to converge towards its asymptotic value as $N^{-2/3}$, exactly as for the SK model, although, as already noted above, this asymptotic value is $\approx 8\%$ below the optimal SK value. Is it possible to characterize more precisely the configurations reached after learning in the long memory limit $\alpha \rightarrow 0$? Can one, in particular, understand analytically the distribution of individual rewards shown in Fig. 5.5 and the corresponding asymptotic value of the average reward, as well as its non monotonic behavior as a function of the noise parameter β ? These are, in our opinion, quite interesting theoretical questions.

6.6.3 Extensions and final remarks

Many extensions of the very simple framework presented here can be imagined for the model to be more representative of real socioeconomic systems. For example, by analogy with spin-glasses, going beyond the fully connected interactions and towards a more realistic network structure should not change the overall phenomenology although some subtle differences may show up. While analytical predictions become even more challenging, recent works on dynamical mean-field theories with finite connectivity, so far developed for Lotka-Volterra type systems, could perhaps be adapted to the learning dynamics of our model [234]. Allowing the interaction network to evolve with time would of course also make the model more realistic, as in [176]; in this case one would have to distinguish the case where the learning time is much longer or much shorter than the reshuffling time of the network.

Other interesting additions to Ising games could also include the introduction of self-excitation [235] or of alternative decision rules that might be less statistical mechanics-friendly [15]. Extension to multinary decisions, beyond the binary case

considered here, as well as higher-order interactions (i.e. a “ p -spin game”), would obviously be interesting as well, specially as higher order interactions are known to change the phenomenology of the SK model (see e.g. [121]). Finally, whereas temperature in physics is the same for all spins, there is no reason to believe that the amount of noise β or the memory span α should be the same for all agents. Introducing such additional heterogeneities might be worth exploring.

Beyond the socioeconomic context that was our initial motivation for its design, we believe that the simplicity and generality of our model makes it a suitable candidate to describe a much wider range of complex systems. In the context of biological neural networks, the parameter α indeed allows one to interpolate between simple discrete-time Hopfield network [35], and continuous-time models where Q_i is an activation variable for the firing rates m_i [236–241]. Although in our case the influence of β introduces some perhaps unwanted stochasticity, these fluctuations can in principle be suppressed (at least partially) with sufficiently small α . The memory loss parameter could also represent an interesting way to tune the effective slowing down of the dynamics caused by symmetry and described in [227]. Here, the description of real neural networks would likely require much more sparse interactions, but also perhaps the introduction of some dedicated dynamics for the interactions themselves, see e.g. [242] for recent ideas.

Closer to our original motivation, more applied socioeconomic problems might benefit from the introduction of this type of reinforcement learning. In macroeconomics for instance, some form of “habit formation” could perhaps be relevant to extend existing descriptions of input/output networks [150, 176], where client/supplier relationships are probably strongly affected by history (on this point, see Kirman’s classic study on Marseilles’ fish market [243]). Finally, while it is an aspect of our model we have not investigated here, previous works have reported that similar dynamics yield interesting volatility clusters and heavy tails that might be interesting to describe financial time series [177, 192].

Last but not least, we believe that the learning dynamics presented here may be useful from a purely algorithmic point of view in the study of spin-glasses and so-called TAP states. Indeed, in the $\beta \rightarrow \infty$ limit, we have seen that our iteration relatively frequently finds fixed points in regions where their abundance is known to be sub-exponential (close to $\varepsilon = 1$ in particular), and this even for relatively large values of α . Interestingly, similar exponentially weighted moving averages have been employed in past numerical studies of TAP states for symmetric interactions [244, 245], but on the magnetizations m_i themselves and not on the local fields Q_i like is the case above.


Using an offline version of our learning procedure could then be of use to effectively converge to fixed points of the TAP equations or Naive Mean-Field Equations, and study their properties. Perhaps even more interestingly, the online dynamics and the resulting fluctuations of the m_i themselves could perhaps be

valuable to probe hardly accessible regions of the solution space. In some sense, the fluctuations related to finite values of α and β could allow define “meta-TAP” states, i.e. closely related TAP states mutually accessible thanks to such extra fluctuations, in the spirit of Langevin dynamics in an energy landscape.

Finally, as mentioned in Sec. 5.3.7 in the zero temperature limit and for $\varepsilon = 1$, it has recently been reported that for a certain range of self-interaction strengths J_d , there appears an exponential number of accessible solutions to the TAP equations that are seemingly not reachable with standard Hopfield dynamics [210]. Preliminary numerical experiments seem to suggest that our learning dynamics find such fixed points, as suggested by their effectiveness for $\varepsilon = 1$ without any form of self-interaction. Beyond existing interest around neural networks, it is clear that such a self-reinforcement, “habit formation” term could also be interesting to study from the socioeconomic perspective [173, 175].

Key takeaways

- The SK-game is difficult to describe analytically.
- Studying the fixed point and limit cycle complexities sheds light on the different regimes one can expect, but the outcome of the $N \rightarrow \infty$ dynamics turns out to be markedly different for $\varepsilon \gtrsim 0.5$.
- These $N \rightarrow \infty$ dynamics can be described with DMFT, which can be interpreted as a sophisticated “representative agent” set of equation.
- In the vicinity of the chaotic region, learning appears to significantly increase the likelihood for finite N dynamics to reach a fixed point.
- In the oscillatory regime, the period of (quasi) limit cycles stretches anomalously as $\alpha^{-1/2}$. These cycles are very fragile with respect to the fluctuations caused by bounded rationality.
- DMFT formalizes the previously mentioned non-stationarity of intentions: whenever agents are sufficiently rational to become decided, the fixed point solution is invalid, even when interactions are reciprocal.
- The boundaries of the non-stationary regime where aging occurs in parameter space and the precise role of the non-reciprocity ε , in particular, remain open questions.



Part IV

Detailed balance violation in the absence of disorder

Chapter 7

A Sakoda-Schelling model with no neighborhoods

I want the hood... to want me back.

Kendrick Lamar

This chapter is based on joint work with Ruben Zakine and Antoine-Cyrus Becharat.

In the previous part of the manuscript, the SK-game was centered around the idea of radical complexity, and the dynamics were *de facto* nonrelaxational due to the non-reciprocity of the random interactions in between agents. While we saw that this gave rise to an extremely rich phenomenology that was studied through the lens of disordered systems, we did not dwell on the decision-making rule, focusing instead on the effect of learning. In this part of the manuscript, we leave the world of spin-glasses to isolate the role of detailed balance violation without disorder, first in an interacting agent-based setting and, in the next chapter, in a much simplified single-body problem.

In this chapter, we first focus on an extension of the Sakoda-Schelling model mentioned in the Chap. 1. This spatially grounded occupation model, in which heterogeneity will only stem from the local nature of the agents' perception, allows us to go back to the impossibility to describe a system with a global free energy discussed in Chap. 4, and to explore the robustness of the phenomenology to different decision rules beyond the classical logit choice.

This chapter is largely taken from [5], which was co-written with R. Zakine under the supervision of M. Benzaquen. Subsections mainly due to the work of R. Zakine and A.-C. Becharat have been removed, while the derivation of the local stochastic evolution equation (not included in the original publication) has been added.

7.1 Motivation

As described in the introduction chapter, both Sakoda [67] and Schelling [69] proposed simple lattice models of idealized cities in hope to understand some aspects of urban segregation in post-WWII American cities, and more widely to describe urban and social dynamics. Each site, representing an accommodation, can be empty or occupied by an agent belonging to one of two sub-populations in the system.

Along the years, these Sakoda-Schelling models have attracted further attention from statistical physicists [246–249], due to their simple microscopic rules and paradoxical macroscopic consequences. To the usual end of bridging the gap from *micro to macro*, mappings onto equilibrium systems were suggested [250], but with limited analytical results.

To gain a more in-depth understanding of the mechanism through which individual choices may lead to sub-optimal collective outcomes, Grauwin *et al.* introduced the modified version of the Schelling model mentioned in the first chapter, with a single type of agent occupying a lattice divided in pre-defined neighborhoods, or blocks [14]. In this occupation model, the agents now base their decisions on the neighborhood density, which is identical for all the agents in a given block. This fixed neighborhood structure then allows to describe analytically the steady-state as the minimizer of a free energy, and to recover a nontrivial phase with suboptimal jam-packed neighborhoods. Thanks to the logit rule, a “linkage function” relating individual moves to a change in a global energy-like function can indeed be found following the lines of Sec. 4.5.1, directly leading to the usual statistical mechanics toolbox. Subsequent works have then explored variations of these different models focusing on the effect of altruistic agents [72], dynamics close to criticality [251–253] or habit formation [254].

Even in the seemingly simpler occupation problem of Grauwin *et al.* [14], several questions persist, both from the socioeconomic and statistical physics perspectives. In particular, the role of the specific decision rule and the precise nature of neighborhoods on the phenomenology of the model remain unclear. Indeed, to allow for the standard techniques of statistical mechanics to be applicable, the choice of the neighborhoods and the dynamics is very constrained. As will be discussed in detail, most non-trivial decision rules lead the system out of thermodynamic equilibrium, requiring calculations that are not always readily tractable. As it is extremely difficult to empirically determine how economic agents actually make decisions, the physics-inspired theoretical analysis of toy models has a significant part to play, in particular to determine the robustness of qualitative findings to specific modeling choices. Besides, as argued earlier in the context of Slutsky matrices, the intrinsically individualistic nature of agent-specific moves in socioeconomic models means that the description of collective behaviors as the

minimization of some global energy is often not possible. Understanding simple out-of-equilibrium dynamics as those that arise from the decision rules presented here is therefore also necessary from the methodological point of view.

The goal of this chapter is therefore to assess, within a general Sakoda-Schelling like occupation model, whether and how the sub-optimal concentration of agents in overly dense regions still occurs out of equilibrium. Most importantly, we relax the assumption of taking a specific decision rule, and no longer require pre-defined block neighborhoods as in [14].

7.2 Presentation of the model

Similar to the original works of Sakoda and Schelling, we consider a city structured as a two-dimensional rectangular lattice composed of $M = L_x \times L_y$ sites (or houses). Each site can be occupied by at most one of the $N(\leq M)$ agents living in this city. On each site of coordinate $\mathbf{r} = (i, j)$, the occupation field n takes the value $n(\mathbf{r}) = 1$ if the site is occupied, $n(\mathbf{r}) = 0$ if it is vacant. It is assumed that each agent k wants to maximize their own utility u_k , which depends on the local density of agents around them. Typically and as in [14], it is natural to think that people like to gather in relatively dense areas to benefit from the city life, but not too dense as crowding might degrade the quality of life. Agents estimate the local density by averaging the occupation field with a probability-density-function kernel G_σ , where σ stands for the interaction range. The kernel is assumed to be isotropic and identical for all agents. The smoothed occupation field $\tilde{n}$ at site $\mathbf{r}$ is thus given by the discrete convolution

$$\tilde{n}(\mathbf{r}) = \sum_{\mathbf{r}'} G_\sigma(\mathbf{r} - \mathbf{r}') n(\mathbf{r}'). \quad (7.1)$$

At each time step, an agent k can decide to move out from their occupied site $\mathbf{r}_k$ and to settle on a new empty site $\mathbf{r}'_k$ where the utility $u[\tilde{n}(\mathbf{r}'_k)]$ might exceed their previous utility $u[\tilde{n}(\mathbf{r}_k)]$.

Once again, we will not take agents to be strict utility maximizers by default. To remain as general as possible, however, the decision rule is simply assumed to be probabilistic, such that

$$\mathbb{P}(\mathbf{r}_k \rightarrow \mathbf{r}'_k) = f_\beta(\Delta u_k), \quad (7.2)$$

where the function f_β is larger than $\frac{1}{2}$ whenever $\Delta u_k > 0$. Naturally, f_β is a positive and monotonic function of the utility difference, with $\lim_{x \rightarrow -\infty} f_\beta(x) = 0$ and $\lim_{x \rightarrow +\infty} f_\beta(x) = 1$. The previously introduced and discussed logit rule is then a specific instance of this function, for which

$$f_\beta(x) = \frac{1}{1 + e^{-\beta x}}. \quad (7.3)$$




Figure 7.1: (a) Utility function $u(\rho) = -|\rho - \rho^*|^\alpha$ for $\rho^* = 0.5$ and $\alpha = \{0.5, 1, 2\}$. Panels (b), (c) and (d) show snapshots of the stationary state for systems of size $L_x = L_y = 100$, for different utility functions, starting from the same homogeneous profile at $\rho_0 = 0.5$. The stationary density ρ_d in the dense phase is $\rho_d = 0.575(5)$ for $\alpha = 0.5$ in (b), $\rho_d = 0.575(5)$ for $\alpha = 1$ in (c), and $\rho_d = 0.585(5)$ for $\alpha = 2$ in (d). These bulk densities are all significantly higher than the density ρ^* for which agents maximize their utility.

The last ingredient to specify is the utility function u of the agents. As stated above, we assume that the utility depends on the locally smoothed occupation $\tilde{n}$ only, and that it is non monotonic. As in Ref. [14], we assume that the utility is maximal for some density $\rho^* \geq \frac{1}{2}$. We specifically choose

$$u(x) = -|x - \rho^*|^\alpha, \quad (7.4)$$

with $\alpha > 0$, see Fig. 7.1(a).

7.3 In or out of equilibrium?

As discussed in Sec. 4.5.1, an often unspoken motivation for the use of the logit rule in the modeling of socioeconomic systems is that it may satisfy detailed balance. Indeed, we have seen that if one manages to find a system-wide energy-like function $\mathcal{H}$ such that

$$\Delta u_k = \mathcal{H}([\{n(\mathbf{r})\}, n(\mathbf{r}'_k) = 1, n(\mathbf{r}_k) = 0]) - \mathcal{H}([\{n(\mathbf{r})\}, n(\mathbf{r}'_k) = 0, n(\mathbf{r}_k) = 1]), \quad (7.5)$$

then the usual tools of equilibrium statistical mechanics can be used (Chap. 2). The steady-state distribution of agents is notably identified as the minimum of the free energy, which is a Lyapunov function of the dynamics prescribed by the logit rule.

At the agent level, the existence of such a global quantity is usually the symptom of either altruistic individuals (that voluntarily maximize some collective satisfaction) or of a central planner (that constructs individual rewards towards a collective objective). Outside of these two cases, we have established that the existence of a free energy when agents are individualistic is in fact restricted to a limited number of carefully chosen models. In the literature of Schelling-like

models, taking a city divided in neighborhoods or blocks [14], where agents share the same utility, yields such a free energy description (which is importantly not a simple aggregation of individual utilities, see the discussion on their model in Chap. 1). In our model, however, this is no longer true.

To explicitly show that the dynamics breaks detailed balance even when taking the logit rule, one may consider a small system and find a specific cycle breaking Kolmogorov's criterion [255]. Such a cycle between four consecutive states with $N = 3$ agents placed on a one-dimensional “city” with $M = 5$ sites is illustrated in Fig. 7.2. The ratio of transition rates between states X and Y , that differ by an agent located on sites i in X , versus j in state Y , is given by

$$\frac{W_{X \rightarrow Y}}{W_{Y \rightarrow X}} = \frac{1 + e^{-\beta[u(\tilde{n}_i^X) - u(\tilde{n}_j^Y)]}}{1 + e^{-\beta[u(\tilde{n}_j^Y) - u(\tilde{n}_i^X)]}} = e^{\beta[u(\tilde{n}_j^Y) - u(\tilde{n}_i^X)]}. \quad (7.6)$$

As a result, the ratio between the product of forward rates, W_+ , and the product of backwards rates, W_- , in the cycle shown in Fig. 7.2, is given by

$$\frac{W_+}{W_-} = e^{\beta[u(\tilde{n}_5^B) - u(\tilde{n}_3^A) - u(\tilde{n}_2^B) + u(\tilde{n}_3^D) + u(\tilde{n}_2^A) - u(\tilde{n}_5^D)]}. \quad (7.7)$$

For a generic non-linear utility function, $W_+ \neq W_-$, which is a signature of nonequilibrium dynamics. For a linear utility function on the other hand, considering that the convolution kernel G_σ is isotropic, all terms in the exponential cancel out, leading to $W_+ = W_-$ (which would be also satisfied for any other cycle). In this situation, the utility difference can simply be interpreted as an energy difference, where the kernel G_σ plays the role of a pairwise interaction potential between the agents. Interestingly, this small cycle also illustrates how the introduction of neighborhoods can salvage the equilibrium description for a generic utility. Splitting the lattice in two neighborhoods along the dashed line shown in Fig. 7.2 and taking an identical value of $\tilde{n}$ for all agents on each neighborhood, the terms in the exponential in Eq. (7.7) indeed cancel out for any utility function since $\tilde{n}_5^B = \tilde{n}_5^D$, $\tilde{n}_3^A = \tilde{n}_2^A$ and $\tilde{n}_2^B = \tilde{n}_3^D$.

7.4 Numerical experiments

Having established the out-of-equilibrium nature of our model, we start by performing numerical simulations to assess whether the concentration of agents in overly dense regions is generic and robust to different shapes of the utility function. Here, all numerical simulations are performed on a two-dimensional grid with periodic boundary conditions. The utility is maximal for $\rho^* = 1/2$. For the sake of simplicity, here we use the logit decision rule and a truncated Gaussian




Figure 7.2: Loop of four configurations with $N = 3$ agents on $M = 5$ sites breaking Kolmogorov's criterion when the utility is non-linear and is a function of an individual perceived density. Shaded and unshaded nodes correspond to occupied and empty sites respectively. The dashed line indicates a possible segmentation of the system into two distinct neighborhoods.

kernel

$$G_\sigma(\mathbf{r}) = \begin{cases} \frac{1}{N_\sigma} e^{-\frac{1}{2\sigma^2} \|\mathbf{r}\|^2}, & \text{if } \|\mathbf{r}\| \leq 4\sigma, \\ 0, & \text{otherwise,} \end{cases} \quad (7.8)$$

where N_σ enforces the normalization of the kernel.

Unless indicated otherwise, all of the numerical data presented in the figures below have been produced by R. Zakine.

7.4.1 Phase separation

For large system size $L_x, L_y \gg \sigma$, we explore the behavior for different global densities $\rho_0 = N/(L_x L_y)$ and for various rationality parameters β . Numerical results are qualitatively similar for all the values of α we tested, ranging from $\alpha = 0.5$ to $\alpha = 2$, see Fig. 7.1. The phenomenology can be summarized as follows. When rationality is low ($\beta \rightarrow 0$, $T \rightarrow \infty$), the stationary state remains homogeneous because agents settle at random. When rationality is high, agents may aggregate in dense clusters, which can surprisingly be more crowded than what agents' utilities prescribe. This was already discussed in [14] where the authors point out that the homogeneous state is actually an unstable Nash equilibrium, even though all agents maximize their utility. The destabilization occurs as one agent randomly moves to another region (with no regard to the effect it may have on the other agents utilities), which decreases the average density at their original site and increases the average density where they settle. Agents in the lower-density region will eventually move to gain the utility they lost when their neighbors moved out. This dynamics will eventually empty some regions, in which agent's return becomes statistically less and less probable. The final state, where a dense phase and an empty phase coexist, is a stable Nash equilibrium.




Figure 7.3: (a) Typical dense domain size $L_d(t)$ during coarsening as a function of time t . A unit of time is defined as N Monte Carlo steps, where N is the number of agents. $L_d(t)$ is averaged over 5 independent simulations. (b), (c), (d) and (e) show snapshots at different times. Starting from a disordered configuration, we quench the system at low temperature, or high rationality β , corresponding to $T \simeq T_c/6$. Parameters: $L_x = L_y = 600$, $\rho_0 = 0.3$, $\sigma = 1$, $\alpha = 3/2$, $T = 0.01$.

One can quantify the condensation dynamics when starting from the homogeneous state and taking high rationality. The system undergoes a spinodal decomposition where dense clusters grow and merge until there is one large dense cluster only, as shown in Fig. 7.3. The final cluster topology ultimately depends on noise realization and on the box dimensions. We measure the cluster size $L_d(t)$ as a function of time t using the radial structure factor. We find $L_d(t) \sim t^{1/z}$, with the dynamical exponent $z \in [2, 3]$, reminiscent of the coarsening exponent observed in a $2d$ Ising system with long-range Kawasaki dynamics [256–259]. Interestingly, and consistent with the findings of [259] in the low temperature region, our results suggest an exponent closer to the local Kawasaki dynamics result $z = 3$ (see Fig. 7.3(a)), despite long-range particle displacements.

7.4.2 Critical point and exponents

The complete phase separation that occurs when rationality is high suggests the use of the order parameter $m \equiv \rho_d - \rho_g$, where ρ_d, ρ_g are the average densities of the dense and “gas” (dilute) phases, respectively. At the critical point (ρ_c, T_c) , we expect a second-order phase transition where m goes to 0 with power-law scaling

$$m \underset{\tau \rightarrow 0^+}{\sim} \tau^\beta, \quad (7.9)$$

where $\tau = (T_c - T)/T_c > 0$ defines the rescaled temperature difference, and β is the order-parameter critical exponent³³. Measuring the critical exponents allows one to determine to which universality class the system belongs to, providing precious

³³Again, this choice of letter seems quite unfortunate given the presence of the rationality/inverse temperature β , even though the typesetting has been adapted.




Figure 7.4: Numerical experiments for $\sigma = 1$, $\alpha = 3/2$. (a) Binodal densities measured for $L_x = 200$ and $L_y = 66$ ($\ell = 33$), inset showing the Binder cumulant as a function of the density and fitted (continuous line) to determine the critical density. (b), (c) and (d) show the numerical measurements of the critical exponents close to the critical point $(\rho_c, T_c) = (0.271, 0.0620)$ determined using various system sizes ranging from $\ell = 20$ to $\ell = 40$.

information on the system behavior at large scales. Since simulations are carried out in finite systems, measuring the critical point with precision requires numerical tricks. We follow the approach that has been extensively used to measure critical exponents in systems undergoing a Motility-Induced Phase Separation (referred to as MIPS) [260–262].

As detailed in [5], simulations are performed in a rectangular domain of size $L_x \times L_y$, with $L_x = 3L_y$, with periodic boundary conditions to keep flat interfaces between a stripe of liquid and a stripe of gas. Four square boxes of sides $\ell = L_y/2$ are then taken, two in the bulk of the dense phase and two in the bulk of the dilute phase. The critical density and temperature can then be identified with the celebrated Binder cumulant, which is a measure of the local density fluctuations. The exponent β is directly measured from the order parameter m as function of reduced temperature τ , at a fixed system size, see Fig. 7.4(c). Other critical exponents can also be measured. By considering the scaling of the Binder cumulant with the system size as shown in Fig. 7.4(b), one can obtain ν , the critical exponent associated to the correlation length. The susceptibility exponent γ can on the other hand be measured directly as was the case with β (Fig. 7.4(d)).

Taking for example $\sigma = 1$ and $\alpha = 3/2$, one identifies the critical point at $\rho_c = 0.271(5)$ and $T_c = 0.0620(2)$, where the uncertainty on the last digit appears in the parentheses. The phase diagram in space (ρ, T) is shown in Fig. 7.4(a), the black star indicating the critical point and the circular markers showing the densities of the coexisting phases: they define the *binodal* frontier. Regarding critical exponents, we measured $\beta \approx 0.12$, $\gamma \approx 1.75$ and $\nu \approx 1.00$. We immediately remark that these values are very close or identical (up to measurement uncertainties) to those associated to the $2d$ -Ising universality class, for which $\beta = 1/8$, $\gamma = 7/4$ and $\nu = 1$. These results enjoin us to assert with a high degree of confidence that the model considered here belongs to the $2d$ -Ising universality class. Since the

system is out of equilibrium and particle displacements can be of infinite range, recovering the Ising universality class is nontrivial and is a remarkable feature, in our opinion. This result must also be put into perspective with the recent debate on the universality class of systems undergoing MIPS [260–263]. Indeed, the interaction kernel G_σ provides a so-called *quorum-sensing* interaction, like that found in assemblies of bacteria [264]. The particle dynamics is however quite different for bacteria and for our agents.

7.5 Mean-field description

We start by writing the expectation of the occupation number $n_{\mathbf{r},s+1} \equiv n(\mathbf{r}, s+1)$ of site $\mathbf{r}$ at time $s+1$, conditioned on the previous configuration $\{n_{\mathbf{r},s}\}$. Averaging over multiple realizations of noise and using a mean-field approximation in which all correlation functions factorize, one obtains

$$\begin{aligned} \langle n_{\mathbf{r},s+1} \rangle - \langle n_{\mathbf{r},s} \rangle &= (1 - \langle n_{\mathbf{r},s} \rangle) \sum_{\mathbf{r}' \neq \mathbf{r}} \langle n_{\mathbf{r}',s} \rangle f_\beta(\Delta u_{\mathbf{r}' \rightarrow \mathbf{r}}^s) \\ &\quad - \langle n_{\mathbf{r},s} \rangle \sum_{\mathbf{r}' \neq \mathbf{r}} (1 - \langle n_{\mathbf{r}',s} \rangle) f_\beta(\Delta u_{\mathbf{r} \rightarrow \mathbf{r}'}^s), \end{aligned} \quad (7.10)$$

where $\Delta u_{\mathbf{r} \rightarrow \mathbf{r}'}^s \equiv u(\langle \tilde{n}_{\mathbf{r}',s} \rangle) - u(\langle \tilde{n}_{\mathbf{r},s} \rangle)$. For convenience, we take the continuous time and continuous space limit, following the common procedure to obtain a mean-field description of exclusion processes on lattices (see e.g. [265]). The average occupation number $\langle n \rangle$ is now described by the density ρ , while the spatially smoothed average occupation number $\langle \tilde{n} \rangle$ is described by the field $\phi \equiv G_\sigma * \rho$. The master equation for the occupation probability then takes the form of a noiseless hydrodynamic equation, in our case:

$$\begin{aligned} \partial_t \rho(x, t) &= [1 - \rho(x, t)] \int dy \rho(y, t) w_\beta([\phi], y, x, t) \\ &\quad - \rho(x, t) \int dy [1 - \rho(y, t)] w_\beta([\phi], x, y, t), \end{aligned} \quad (7.11)$$

with the transition rate from y to x explicitly given by

$$w_\beta([\phi], y, x, t) = \omega f_\beta [u(\phi(x, t)) - u(\phi(y, t))], \quad (7.12)$$

where ω is homogeneous to an inverse time scale. Eq. (7.11) is valid in any dimension, but, for simplicity, we will work out the mean-field computations in dimension 1 in space. This can be justified *a posteriori* when we compare the mean-field (MF) to the Monte Carlo (MC) simulations. Let us also mention that

the dimension does not play a role in determining the phase densities in the steady-state of coarse-grained field theories (Allen-Cahn [266], Cahn-Hilliard [267], etc.).

Integrating Eq. (7.11) over space, one immediately sees that the total density $\int \rho$ is conserved. One can also check that in the very specific case where $u(\phi)$ is linear in ϕ , one can build a free energy functional that is a Lyapunov function of the non-local MF dynamics, ensuring a convergence towards local minima and preventing limit cycles and oscillatory dynamics. This is a natural consequence of the fact that detailed balance is satisfied at the microscopic level. In App. D.1, we construct this free energy and show that the dynamics are relaxational.

7.5.1 Linear stability analysis

In the general case, we would like to understand how the homogeneous state becomes unstable. To do so, we consider a small perturbation around the homogeneous state: $\rho(x, t) = \rho_0 + \rho_1(x, t)$, with ρ_1 the perturbation. By linearity of the convolution, one has $\phi(x, t) = \rho_0 + \phi_1(x, t)$, with $\phi_1 \equiv G_\sigma * \rho_1$. A Taylor expansion of Eq. (7.11) combined with mass conservation (i.e. $\int_D \rho_1 = \int_D \phi_1 = 0$, where D is the full domain), finally yields:

$$\partial_t \rho(x, t) = 2\Omega \rho_0 (1 - \rho_0) f'_\beta(0) u'(\rho_0) \phi_1(x, t) - \Omega f_\beta(0) \rho_1(x, t), \quad (7.13)$$

with Ω the full domain size. Defining the Fourier transform for any field h as $\hat{h}(k) = \int dx e^{-ikx} h(x)$, one obtains

$$\partial_t \hat{\rho}_1(k, t) = \Lambda(k) \hat{\rho}_1(k, t), \quad (7.14)$$

$$\Lambda(k) = \Omega f_\beta(0) \left(2\rho_0 (1 - \rho_0) \frac{f'_\beta(0)}{f_\beta(0)} u'(\rho_0) \hat{G}_\sigma(k) - 1 \right). \quad (7.15)$$

This last equation shows that the homogeneous state is unstable if there exists a mode k^* such that

$$2\rho_0 (1 - \rho_0) \frac{f'_\beta(0)}{f_\beta(0)} u'(\rho_0) \hat{G}_\sigma(k^*) > 1. \quad (7.16)$$

The manifold for which the inequality becomes an equality defines the spinodal in the phase diagram (ρ_0, β) . In particular, for any monotonically decreasing kernel $G_\sigma(|x|) \in L_2(\mathbb{R})$, one has $\hat{G}_\sigma(0) > |\hat{G}_\sigma(k)|$, such that for large system size, the stability of the homogeneous state is given by the stability of modes $k \rightarrow 0$, and the spinodal is thus defined by the equation

$$2\rho_0 (1 - \rho_0) \frac{f'_\beta(0)}{f_\beta(0)} u'(\rho_0) = 1. \quad (7.17)$$

Note that this criterion is generic as it only depends on the decision rule through $f_\beta(0)$ and $f'_\beta(0)$. The simulations also reveal the existence of a bistable region in the vicinity of this spinodal. This is the binodal region, which can be fully characterized in the case of an equilibrium system [248]. Here however, there is a priori no free energy one can rely on to describe the nucleation scenario and to obtain the densities of the phase-separated state.

7.5.2 Comparison with numerical simulations

The MF prediction is expected to be accurate for systems with high connectivity, which here corresponds to large σ . In the following, we shall take the limit $L \rightarrow \infty$, $\sigma \rightarrow \infty$ with $\sigma/L \rightarrow 0$ to obtain mean-field predictions that are independent of both σ and L , and perform numerical simulations as close as possible to this scaling regime.

The first analytical prediction of the MF description is the spinodal, that determines the onset of instability of the homogeneous state, see Eq. (7.17). The spinodal is the dashed line in the (ρ, T) phase diagram in Fig. 7.5(a). To check the prediction, we start in the MC simulations from a uniformly distributed configuration of agents for three different values of temperature, $T = 0.04, 0.08, 0.11$, and we detect the frontier across which the homogeneous profile either coarsens, or needs a nucleation event to converge to the separated state. This frontier is marked with the diamonds, which agrees with the MF prediction.

Second, the MF dynamical Eq. (7.11) can be solved numerically with an Euler explicit scheme. From the numerical solution, one obtains the densities of the bulk of each phase when a phase separation occurs: these densities define the binodal, the continuous line in Fig. 7.5(a). These MF phase densities are perfectly recovered by the MC simulations (circles). In addition, one can compare the steady-state average density profile from MC simulations to the mean-field stationary density, which superimpose almost exactly, see Fig. 7.5(b).

As previously stated, the MF predictions fail for small values of σ . The phase diagram in Fig. 7.4(a) is for instance obtained for $\sigma = 1$, and indeed strongly differs from the MF solution. For $\sigma = 1$, we notably identified the critical point as $(\rho_c, T_c) = (0.271, 0.0620)$, whereas the MF predicts $(\rho_c, T_c)_{\text{MF}} = (0.2763, 0.1418)$, where, as expected, $T_c^{\sigma=1} < T_c^{\text{MF}}$.

7.5.3 Local move approximation

To make progress into the identification of a possible effective free energy functional, it may be convenient to consider slightly modified dynamics where jumps are now only authorized in the direct neighborhood of the agents. Indeed, considering an evolution enforcing a *local* mass conservation will allow for more familiar




Figure 7.5: (a) Phase diagram for $\alpha = 3/2$, $\sigma = 7$ and $L_x = 200$, $L_y = 66$ ($\ell = 33$). The dot-dash black line corresponds to the binodal densities of the numerical solution to the mean-field Eq. (7.11), circles are measured from Monte Carlo (MC) simulations. The dashed black line represents the mean-field spinodal from linear stability analysis, while diamonds indicate the loss of stability of the homogeneous state in the MC simulations. The black square is the critical point for $\sigma/L \rightarrow 0$. (b) Averaged density profile $\rho(x)$ from MC simulations (continuous green line) for the same parameters as (a) and $\rho_0 = 0.35$, $T = 0.05$. The dashed black line is the stationary numerical solution of the mean-field equation Eq. (7.11).

partial differential equations (PDEs) and field theoretic approaches on conserved scalar fields. Here, the absence of macroscopic density currents in the steady-state, both in MC simulations and in the MF solution suggests that the system generically converges to a stationary stable fixed point, where the details of the dynamics become inconsequential. In addition, when the majority of agents have aggregated in a single dense cluster in the steady-state, it is unlikely that they would perform moves outside of the bulk, in low-density regions, since the utility there is minimal. The local-move approximation, as it strongly simplifies the description, thus appears natural.³⁴

Following the Taylor expansion outlined in App. D.2, the local mean-field dynamics is given by

$$\partial_t \rho = f_\beta(0) \partial_x^2 \rho - 2f'_\beta(0) \partial_x [\rho(1 - \rho) \partial_x u], \quad (7.18)$$

which can be rewritten as the canonical equation for the mass-conserving dynamics

$$\partial_t \rho = \partial_x [M[\rho] \partial_x \mu([\rho], x)], \quad (7.19)$$

with the mobility operator $M[\rho] = \rho(1 - \rho)$, stemming from the non-overlapping

³⁴Dynamically, the coarsening exponent $z \simeq 3$ displayed in Fig. 7.3(a), and which is also observed in a Cahn-Hilliard relaxation dynamics can also be invoked to support the idea of local moves.

nature of the agents, and with the chemical potential $\mu = \mu_{\text{ent.}} + \mu_{\text{util.}}$ where

$$\mu_{\text{ent.}} = f_{\beta}(0) \log \left(\frac{\rho}{1 - \rho} \right) \quad (7.20)$$

$$\mu_{\text{util.}} = -2f'_{\beta}(0)u[\phi(x)]. \quad (7.21)$$

The first contribution to the chemical potential $\mu_{\text{ent.}}$ is purely local and accounts for entropy in the system where agents cannot overlap. The second contribution $\mu_{\text{util.}}$ encodes the drive from agents' utility. This term exhibits non-locality with respect to the field ρ , and as a consequence, cannot be expressed as a functional derivative of any free energy, in general [268–270]. However, in the particular case of a linear utility in ϕ , one again recovers that $\mu_{\text{util.}} + \mu_{\text{ent.}}$ can be written as the functional derivative of the free energy $\mathcal{F}$ given in App. D.1. Let us emphasize that, here again, the decision rule is kept general, and that the entire local dynamics only depend on it through $f_{\beta}(0)$ and $f'_{\beta}(0)$.

Performing the linear stability analysis on the dynamics with local moves (see App. D.2), we find that the criterion for the homogeneous solution to be unstable is identical to that given in Eq. (7.17), when moves are global. Also, the stationary density profiles computed either with the local, or with the non-local MF PDEs for the same parameters are identical. Both these observations therefore allow us to confirm the relevance of the local-move approximation to characterize the system in the long-time limit.

7.5.4 Alternative derivation and fluctuating hydrodynamics

Within the local move approximation, we may perform an alternative derivation of the mean-field equation, which will also yield the fluctuations about the deterministic hydrodynamics. Following Biroli & Lefèvre [271], we may write the generating function for particles (locally) diffusing on a lattice as

$$Z[\{n, \hat{n}\}] = \int \{dn d\hat{n}\} e^{-S[\{n, \hat{n}\}]}, \quad (7.22)$$

with the Martin-Siggia-Rose-Jansen-de Dominicis (MSRJD) action

$$S[\{n, \hat{n}\}] = - \int dt \left\{ - \sum_i \hat{n}_i \partial_t n_i + \sum_{(i,j)} n_i W_{ij} (e^{\hat{n}_j - \hat{n}_i} - 1) \right\}, \quad (7.23)$$

where W_{ij} is the transition rate from site i to j . In the case of an exclusion process such as the one we are considering here, we take

$$n_i W_{ij} = n_i (1 - n_j) f_{ij}, \quad (7.24)$$

Chapter 7. A Sakoda-Schelling model with no neighborhoods

such that it takes the value 0 if site i is empty and/or if site j is occupied. Now, we assume that the function f can be locally expanded as $f_{ij} = f(0) + a\mathbf{e}_{ij} \cdot \nabla f(0) + \mathcal{O}(a^2)$, where a is the lattice spacing and $\mathbf{e}_{ij}$ is the unit vector pointing in the correct direction (we are generically considering d dimensional lattices). We can similarly expand n_j but need to go to an additional order for $\hat{n}_j$ in to recover the correct expansion

$$e^{\hat{n}_j - \hat{n}_i} - 1 = a\mathbf{e}_{ij} \cdot \nabla \hat{n}_i + \frac{a^2}{2}[(\mathbf{e}_{ij} \cdot \nabla)^2 \hat{n}_i + (\mathbf{e}_{ij} \cdot \nabla \hat{n}_i)^2]. \quad (7.25)$$

Combining all terms and expanding up to order a^2 , the sum over neighbouring sites in the action gives

$$\begin{aligned} \sum_{(i,j)} n_i W_{ij} (e^{\hat{n}_j - \hat{n}_i} - 1) &= \sum_{(i,j)} \left[af(0)n_i(1-n_i)\mathbf{e}_{ij} \cdot \nabla \hat{n}_i + \frac{a^2}{2}f(0)n_i(1-n_i)(\mathbf{e}_{ij} \cdot \nabla)^2 \hat{n}_i \right. \\ &\quad + \frac{a^2}{2}f(0)n_i(1-n_i)(\mathbf{e}_{ij} \cdot \nabla \hat{n}_i)^2 - a^2 f(0)n_i(\mathbf{e}_{ij} \cdot \nabla n_i)(\mathbf{e}_{ij} \cdot \nabla \hat{n}_i) \\ &\quad \left. + a^2 n_i(1-n_i)(\mathbf{e}_{ij} \cdot \nabla f(0))(\mathbf{e}_{ij} \cdot \nabla \hat{n}_i) \right]. \end{aligned} \quad (7.26)$$

The first term of order a vanishes by symmetry, as the unit vector retains its sign and adjacent sites thus cancel out as $a \rightarrow 0$. In this continuous limit, the sum becomes an integral (time may be rescaled as required), and the action becomes

$$\begin{aligned} S[\{n, \hat{n}\}] &= - \int dt \int d\mathbf{x} \left\{ -\hat{\rho} \partial_t \rho + f(0)\rho(1-\rho)\nabla^2 \hat{\rho} + f(0)\rho(1-\rho)(\nabla \hat{\rho})^2 \right. \\ &\quad \left. - 2f(0)\rho \nabla \rho \cdot \nabla \hat{\rho} + 2\rho(1-\rho)\nabla f(0) \cdot \nabla \hat{\rho} \right\}, \end{aligned} \quad (7.27)$$

where the factor 2 comes from the fact that summing squared terms over all neighboring sites j always gives pairs of terms (for example on a $d = 1$ lattice we have a derivative going forward and backward that are both evaluated at i and so give twice an identical term). At this stage, one can finally perform several integration by parts to recover only terms factorized by $\hat{\rho}$ of $(\nabla \hat{\rho})^2$, corresponding to the deterministic and fluctuating contributions of the final action respectively.

Skiping details of this straightforward procedure, the generating functional is finally given by

$$\begin{aligned} S[\{n, \hat{n}\}] &= - \int dt \int d\mathbf{x} \left\{ -\hat{\rho} \partial_t \rho + f(0)\hat{\rho} \nabla^2 \rho - 2\hat{\rho} \nabla \cdot (\rho(1-\rho)\nabla f(0)) \right. \\ &\quad \left. + f(0)\rho(1-\rho)(\nabla \hat{\rho})^2 \right\}, \end{aligned} \quad (7.28)$$

corresponding to the Langevin equation

$$\partial_t \rho = \nabla \cdot \left[f(0) \nabla \rho - 2\rho(1-\rho) \nabla f(0) + \sqrt{\rho(1-\rho)} \xi \right] \quad (7.29)$$

with the Gaussian white noise ξ correlated as

$$\langle \xi(\mathbf{x}, t) \xi(\mathbf{x}', t') \rangle = 2f(0) \delta^d(\mathbf{x} - \mathbf{x}') \delta(t - t'). \quad (7.30)$$

The deterministic part of the dynamics correctly recovers the previously derived mean-field expression of Eq. (7.19). As such, the fluctuating hydrodynamics in a single dimension can finally be written as before,

$$\partial_t \rho = \partial_x [M[\rho] \partial_x \mu([\rho], x) + \sqrt{2f(0)M[\rho]} \eta], \quad (7.31)$$

with η a Gaussian white noise, and the same chemical potential $\mu = \mu_{\text{ent.}} + \mu_{\text{util.}}$ and mobility operator $M[\rho]$ as before. While not central to the present study, these fluctuations can be studied in more detail, providing information on the nucleation scenarii and on transition paths between macroscopic states for instance [272–275].

Unfortunately, the procedure described here does not appear suited to an extension to non-local moves. Indeed, if i and j are not neighbors in Eq. (7.23), then we can unfortunately not expand the exponential term in $\hat{n}_j - \hat{n}_i$. While not having gradient expansions for all term is not necessarily an issue, it seems that we must be able to have local terms in the conjugate variable to identify the deterministic and fluctuating contributions at the very end.

7.6 Generalized thermodynamic mapping

Even though μ in Eq. (7.19) cannot be written as a functional derivative [268–270], the dynamics can be analyzed by resorting to a gradient expansion. Indeed, expanding the chemical potential up to $O(\nabla^4, \rho^2)$ terms yields

$$\mu[\rho] = g_0(\rho) + \lambda(\rho)(\nabla \rho)^2 - \kappa(\rho)\nabla^2 \rho + O(\nabla^4, \rho^2), \quad (7.32)$$

with g_0 , λ , κ local function of the field ρ , and a generalized thermodynamic mapping [276, 277] can yield the prediction of the binodal densities.

For simplicity, we will now assume that the smoothing kernel is a Gaussian distribution of zero mean and variance σ^2 . In Fourier space, the smoothed field is given by $\hat{\phi}(k) = \hat{\rho}_k \exp(-\sigma^2 k^2/2)$, which can be truncated to leading order:

$$\hat{\phi}_k \simeq \hat{\rho}_k \left(1 - \frac{k^2 \sigma^2}{2} + O(\sigma^4 |k|^4) \right). \quad (7.33)$$

Chapter 7. A Sakoda-Schelling model with no neighborhoods

In real space, this translates into $\phi = \rho + \frac{\sigma^2}{2} \nabla^2 \rho + O(\nabla^4, \rho)$. This allows us to further expand the $\mu_{\text{util.}}$ given in Eq. (7.21). To leading order in the $O(\nabla, \rho)$ expansion, one has

$$\mu_{\text{util.}} = -2f'_\beta(0) \left[u(\rho) + \frac{\sigma^2}{2} u'(\rho) \partial_x^2 \rho + O(\partial_x^4, \rho) \right]. \quad (7.34)$$

Combining this expansion of $\mu_{\text{util.}}$ with the entropic contribution $\mu_{\text{ent.}}$, it is now possible to identify the different terms in Eq. (7.32), namely:

$$g_0(\rho) = -2f'_\beta(0)u(\rho) + f_\beta(0) \log \left(\frac{\rho}{1-\rho} \right); \quad (7.35)$$

$$\lambda(\rho) = 0; \quad \kappa(\rho) = f'_\beta(0)\sigma^2 u'(\rho). \quad (7.36)$$

This identification enables us to follow up to the next step, which is finding the proper function $R(\rho)$ and the generalized functional $\mathcal{G}[R]$ by means of which the dynamics will be given by

$$\partial_t \rho(x, t) = \partial_x \cdot \left[M[\rho] \partial_x \frac{\delta \mathcal{G}}{\delta R(x, t)} \Big|_{R(\rho)} \right], \quad (7.37)$$

allowing the identification of the fixed points through the study of $\mathcal{G}$.

Since $\lambda(\rho) = 0$ in our case, the differential equation that the function R must satisfies (see [276, 277]) is

$$\kappa(\rho) R''(\rho) = -\kappa'(\rho) R'(\rho), \quad (7.38)$$

which simplifies into $(\kappa R')' = 0$, where the $'$ denotes the derivative with respect to ρ . Given R is *a priori* bijective, one can define a new chemical potential $g[R] \equiv \mu[\rho(R)]$. As described in [276, 277], the effective free energy functional $\mathcal{G}[R]$ can then be obtained by requiring

$$g = \frac{\delta \mathcal{G}}{\delta R}, \quad \text{with} \quad \mathcal{G} = \int dx \left[\Phi(R) + \frac{\kappa}{2R'} (\partial_x R)^2 \right], \quad (7.39)$$

with the generalized free energy density $\Phi(R)$. We can finally recover the coexistence densities by performing a double tangent construction, or equivalently a Maxwell construction, on Φ and tracking back to ρ thanks to the bijectivity of $R(\rho)$. For more details on the procedure, as well as the challenges posed by the change of sign in κ caused by the non-monotonous nature of the utility function, see the dedicated section, mainly due to R. Zakine, in [5].

An illustration of the change of variable to $R(\rho)$ is shown in Fig. 7.6(a), while a specific example of the effective free energy density Φ is visible in the inset of Fig. 7.6(b). Numerically implementing the Maxwell construction to efficiently




Figure 7.6: (a) Utility $u(\rho)$ (solid line), and thermodynamic mapping change of variable $R(\rho)$ (dashed line) for $\alpha = 3/2$ and $\rho^* = 1/2$. (b) Comparison between the semi-analytical prediction (dark line) and the binodal densities both obtained via the MC simulations (green circles) and solving numerically the mean-field Eq. (7.11) (dot-dash green line). Inset: double-tangent construction on $\Phi(R)$ for $T = 0.083$.

explore a range of temperatures yields a semi-analytical solution for the binodal, which is compared to the MC results in Fig. 7.6(b), displaying an excellent match. Note that the interaction range σ , that appears in κ but not in Φ , does not play a role in the predicted coexistence densities of the infinite size system, thereby confirming that the sub-optimal aggregation of agents in a dense cluster is not limited to finite size lattices.

7.7 Two population extension

A natural extension of the problem is to restore some diversity among agents, as initially considered by both Sakoda and Schelling. Here we consider two types of interacting agents (say A and B), with possibly different utility functions, which could for example represent higher and lower revenue individuals, or city dwellers and business professionals, etc. A central question in this case is whether the system reaches fixed points, or if more complicated dynamics can persist in the long time limit, especially if the two populations have competing interests. Recent work has been devoted to studying nonreciprocal interactions between different kinds of particles, exhibiting the wealth of possible dynamical behavior when particle displacements are local [278, 279]. An interesting question in our setup is for instance: do propagating waves (or frustrated states) survive when nonlocal moves are allowed? Indeed, one may expect that enforcing local displacement constitutes a dynamical constraint that drives the system in a particular way. Allowing for nonlocal moves may change the dynamics of how the frustrated states are resolved. One may think of three major types of interactions:

- First, a cooperative interaction where agents A and agents B may maximize

their utility when agents of opposite type are found in their neighborhood. This kind of interaction will typically lead to homogeneous well-mixed systems, or to some condensation into a dense phase where agents are well-mixed, but since frustration is not implemented in the microscopic rules, we reasonably expect stationary states.

- Second, each agent type may decide to settle among peers and/or avoid agents of the other type in their surroundings. One should then expect a complete phase separation into two domains, one displaying a majority of A s and, the other, a majority of B s. Whether the $A - B$ phase separation additionally displays some condensation depends on the self-affinity of each agent type.
- Third, frustrated situations in which A settles with A but wants to avoid B agents, while B agents would like to gather and settle close to A . In this situation, we may expect non stationary patterns, stemming from the fact that all agents cannot be satisfied at the same time.

With this last situation in mind, we have considered the following utility functions (u_A for A agents and u_B for B agents):

$$u_A(x, [\phi_{A,B}]) = -|\phi_A(x) - \rho^*|^2 + c_1 \phi_B(x) \quad (7.40)$$

$$u_B(x, [\phi_{A,B}]) = -|\phi_A(x) - \rho^*|^2 + c_2 \phi_B(x), \quad (7.41)$$

where $c_1 < 0$ translates the fact that A s are fleeing from B , and $c_2 > 0$ translates the fact that B s have a tendency to gather with B s. The term $-|\phi_A - \rho^*|^2$ enjoins both populations to settle among A populated areas. Of course, the specific shape of utilities taken here may be restrictive and can be easily generalized.

The extension of the mean-field dynamics to this two population problem is rather straightforward. Writing $\rho_A(x, t)$ (resp. $\rho_B(x, t)$) the density of agents A (resp. B) at location x and time t , and denoting the total density by $\rho(x, t) \equiv \rho_A(x, t) + \rho_B(x, t)$, we now have an evolution equation of the form

$$\begin{aligned} \partial_t \rho_A(x, t) = & [1 - \rho(x, t)] \int \rho_A(y, t) w_{\beta_A}([\phi_{A,B}], y, x, t) dy \\ & - \rho_A(x, t) \int [1 - \rho(y, t)] w_{\beta_A}([\phi_{A,B}], x, y, t) dy, \end{aligned} \quad (7.42)$$

and, by symmetry, a similar equation for B . The transition rates depend on the utility function of each agent type and are *a priori* agent specific. Denoting $u_Z(x) \equiv u_Z(x, [\phi_{A,B}])$ (with $Z = A$ or B), we set

$$w_{\beta_Z}([\phi_{A,B}], y, x, t) = \omega_Z f_{\beta_Z}[u_Z(x) - u_Z(y)], \quad (7.43)$$




Figure 7.7: Snapshots of the system for two frustrated interaction parameter choices. (a) Stationary demixing in a region where LSA presents complex eigenvalues. The agent A phase still contains some B agents. Parameters: $c_1 = -2$, $c_2 = 1$, $\sigma = 3$, $\bar{\rho}_A = 0.2$, $\bar{\rho}_B = 0.5$, $\beta = 10$. (b) Chaotic propagation of polarized blobs in a region where LSA presents pure real eigenvalues (null imaginary part). Parameters: $c_1 = -2$, $c_2 = 0.5$, $\sigma = 7$, $\bar{\rho}_A = 0.6$, $\bar{\rho}_B = 0.2$, $\beta = 100$. For both (a) and (b), $L_x = L_y = 300$.

where ω_Z and β_Z can be agent dependent.

In App. D.3, we perform the linear stability analysis of the homogeneous state. As expected, in the frustrated two-population system, unstable modes can display temporal oscillations. However, these oscillations may stop when nonlinear terms become relevant, and the system may end up in a stationary phase separation (similar to classical demixing in equilibrium systems), as displayed in Fig. 7.7(a). Reciprocally, non-oscillating growing modes at the linear level may give rise to propagating structures and waves when nonlinearities become important, as shown in Fig. 7.7(b) (see Supplementary Material in [5]). In our system, and at odds with recent work [278, 279], the oscillatory nature of the non-homogeneous steady-state cannot be predicted from a simple linear stability analysis about the homogeneous solution.

A thorough analysis of the emerging behaviors in the multi-population system would require more work, beyond the scope of the present chapter. Still, it is remarkable that, here as well, the linear stability analysis in the case of local jumps yields exactly the same instability conditions as the nonlocal dynamics ones (see results of Appendices D.3 and D.4). As a consequence, we expect that some results of the recent works [278, 279] should be relevant, to some extent, to describe our multi-population system.

7.8 Conclusion

Let us summarize what we have achieved in this chapter. We have introduced a neighborhood-less extension of the Sakoda-Schelling model for the occupation of a lattice representing a city. In this version of the model, the agents attempt to

maximize a utility that is a function of their *perceived* local density, and are most satisfied when they are located in a site where such density is of an intermediate value, i.e. neither empty nor too crowded. Having that agents only consider their own site dependent utility, and that their utility is nonlinear, drives the system out of equilibrium. As a result, the system can no longer be studied by constructing a free energy directly from an aggregate system-wide utility function, as was done by Grauwin *et al.* [14] for agents inhabiting predefined neighborhoods or blocks in which the utility is identical for all (see Sec. 1.3.1). Using numerical simulations as well as a mean-field description of the nonequilibrium dynamics, we have established that the apparent disparity between *micromotives* and *macrobehaviours* initially observed by Schelling [69] is robust to the absence of neighborhoods and to the out-of-equilibrium nature of our extension. Interestingly, we find that the transition between the fully homogeneous state and the phase-separated one likely belongs to the $2d$ Ising universality class, a debated topic in the active matter literature regarding the very similar Motility Induced Phase Separation (MIPS) phenomenon. Taking advantage of the similarity between our problem and the Active Model B (describing MIPS), we predict the local density in the bulk of the concentrated phase, confirming that the majority of agents find themselves in a sub-optimal situation with a perceived density exceeding the ideal value.

While seemingly technical, the fact that the original observations of Schelling is robust to out-of-equilibrium dynamics actually carries far reaching consequences, in our opinion. Indeed, as discussed in Sec. 7.3, equilibrium descriptions of socioeconomic problems require the decision rule $f_\beta(x)$ to be the logit function. This very specific choice is a common source of criticism, as any results are then *a priori* uncertain to hold for other decision rules. Here, on the other hand, our out-of-equilibrium description presents no such restriction, as all calculations have been written as generally as possible and, interestingly, turn out to only depend on $f_\beta(0)$ and $f'_\beta(0)$. While the final results presented here are those using the classical choice of the logit function for lack of a more plausible decision rule, one could readily adapt the outcomes following behavioral evidence that another function is more appropriate, and we expect the results to hold qualitatively for any other sigmoid function. More generally, we believe that this robustness to other decision rules holds for a large number of socioeconomic models that have been described using the methods of statistical physics, for instance that of Chap. 4.³⁵ Of course, subtleties around the dynamics, such as the relaxation time towards the steady-state or the coarsening dynamics discussed here, will inherently be affected by the specific choice that governs transition rates. This being said, we have observed a remarkable similarity in the local and non-local versions of our model


³⁵One might also note that in the different context of the SK game, weakly non-reciprocal interactions (which also violate detailed balance) were also found to leave the phenomenology of the model largely unaffected.

for which the dynamics are yet qualitatively very different. It is therefore possible that there also exists some degree of universality in the dynamical behavior of different decision rules, at least at the mean-field level.

Going back to the Sakoda-Schelling model, we have also considered the introduction of different sub-populations, leading to dynamical patterns that are impossible to observe in an equilibrium version of the model. Another important extension would be to introduce a housing market, and to determine under which conditions such a market might destroy condensation – see the dedicated section presenting the work of A.-C. Becharat in [5]. In order to identify the plausible effects of the housing market, a thorough analysis of real estate transactions appears to be a promising avenue, in particular given the increasing availability of open datasets in this area in major European cities. An extensive study of French data is currently underway, hopefully allowing the coupling of this continuous Sakoda-Schelling model with realistic price dynamics in the near future [280].

Key takeaways

- Replacing fixed neighborhoods by a space-dependent perception of the density drives the Sakoda-Schelling model out of equilibrium, regardless of the decision rule.
- When the perception range is large, the *spinodal* delimiting the region where a homogeneous density is linearly stable can be determined analytically with a mean-field approximation.
- Due to the out-of-equilibrium nature of the system, the coexistence densities in the phase separated region, defining the *binodal*, cannot be determined from a free energy minimization.
- A generalized thermodynamic mapping, originally developed for the theory of active matter, can be employed to circumvent this difficulty.



- We find the sub-optimal aggregation of agents in concentrated clusters to be largely unaffected by the nonequilibrium nature of the model and the specific choice of decision rule.
- The field-theoretic description of the model offers many possible extensions, notably to several sub-populations and to the coupling with a price field.

Chapter 8

Steady-state preserving out-of-equilibrium currents

急がば回れ

More haste, less speed.

Japanese proverb

This chapter is based on joint work with Andreas Dechant.

Throughout this thesis, we have seen how the specificities of socioeconomic systems may lead to out-of-equilibrium dynamics that cannot be described with the equilibrium statistical mechanics toolbox. While we have more or less been able to study these out-of-equilibrium systems with analytical methods, such as DMFT or the generalized thermodynamic mapping, these ultimately required some rather involved numerical schemes and were by no means systematic.

In order to acquire a better understanding of nonequilibrium steady-states (NESS), an interesting approach that has emerged is to purposefully construct out-of-equilibrium models with a *prescribed* steady-state distribution, corresponding to that of a known equilibrium system. As hinted in Chap. 2, it is indeed possible to sample an equilibrium distribution while violating detailed balance. In addition to allowing one to compute many observables thanks to the explicit knowledge (and intuitive understanding) of the steady-state distribution – a rare occurrence out-of-equilibrium –, we will see that these peculiar systems may also provide an appreciable speedup in the sampling of complex equilibrium distributions. Beyond this practical interest, this chapter also allows us to get acquainted with the theory of stochastic thermodynamics, which has brought some significant progress in the understanding of out-of-equilibrium systems in recent years.

8.1 General idea

In the following, we will focus on Langevin dynamics as they allow for a visual interpretation of the out-of-equilibrium currents considered, but most of the ideas and results discussed can be transposed to discrete Markovian dynamics.

8.1.1 Physical setup

In order to derive analytical results, we will turn to the very simple case of a *single* particle under the influence of a fixed force field. The dynamics are assumed to follow the overdamped Langevin equation introduced in Sec. 2.1.3,

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t)) + \sqrt{2T}\boldsymbol{\xi}(t), \quad \mathbf{x} \in \mathbb{R}^d \quad (8.1)$$

where $\boldsymbol{\xi}$ is a vector of independent, identically distributed white noises, i.e. with zero mean and δ -correlated, $\langle \xi_i(t)\xi_j(s) \rangle = \delta_{i,j}\delta(t-s)$. We saw that such dynamics satisfy detailed balance if the force is conservative,

$$\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}), \quad (8.2)$$

where U is the potential (i.e. energy) describing the equilibrium steady-state $P(\mathbf{x}) \propto e^{-\beta U(\mathbf{x})}$, where as before we have $\beta = 1/T$ the inverse temperature.

Suppose we now introduce a detailed balance violating (i.e. irreversible) component $\mathbf{F}_{\text{neq}}(\mathbf{x})$ to the force field,

$$\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}) + \mathbf{F}_{\text{neq}}(\mathbf{x}), \quad (8.3)$$

and recall the *global* balance condition ensuring that $P(\mathbf{x})$ is the steady-state solution to the Fokker-Planck equation,

$$\nabla \cdot [\mathbf{F}(\mathbf{x})P(\mathbf{x}) + T\nabla P(\mathbf{x})] = 0. \quad (8.4)$$

If we require that $P(\mathbf{x})$ remains the Gibbs-Boltzmann distribution of the original equilibrium system, then we necessarily have $-\nabla U(\mathbf{x})P(\mathbf{x}) + T\nabla P(\mathbf{x}) = 0$. As a result, the condition for the out-of-equilibrium forcing to preserve this steady-state distribution is simply

$$\nabla \cdot (\mathbf{F}_{\text{neq}}(\mathbf{x})P(\mathbf{x})) = 0, \quad (8.5)$$

or more explicitly

$$\nabla \cdot \mathbf{F}_{\text{neq}}(\mathbf{x}) - \beta \mathbf{F}_{\text{neq}}(\mathbf{x}) \cdot \nabla U(\mathbf{x}) = 0. \quad (8.6)$$

The resulting out-of-equilibrium system then has a non-zero average steady-state entropy production rate given by

$$\sigma \equiv \frac{1}{T} \langle \mathbf{F}_{\text{neq}}^2 \rangle > 0, \quad (8.7)$$

where angular brackets indicate an average over the steady-state distribution [86]. As a result, we have an irreversible NESS with a known steady-state distribution.

8.1.2 Relaxation to the NESS

Suppose we now have an out-of-equilibrium forcing field satisfying Eq. (8.5) and therefore preserving the steady-state distribution $P(\mathbf{x}) \propto e^{-\beta U(\mathbf{x})}$. While it is clear that a central difference between the original system and its out-of-equilibrium counterpart is the presence of probability currents and therefore a non-zero entropy production rate in the steady-state, the relaxation to this steady-state is also altered.

Rewriting the Fokker-Planck equation in an operator formalism, we have

$$\partial_t P(\mathbf{x}, t) = \mathcal{L}(\mathbf{x})P(\mathbf{x}, t), \quad \mathcal{L}(\mathbf{x}) = -\nabla \cdot \mathbf{F}(\mathbf{x}) + T\nabla^2. \quad (8.8)$$

Now, in the equilibrium case, the forcing field is purely gradient, implying $\partial_{x_j} F_i = \partial_{x_i} F_j$. As a result, one can show that the equilibrium Fokker-Planck operator is self-adjoint,

$$\mathcal{L}_{\text{eq}}(\mathbf{x}) = \mathcal{L}_{\text{eq}}^\dagger(\mathbf{x}), \quad (8.9)$$

which does not hold for out-of-equilibrium dynamics [92].

To solve the time-dependent Fokker-Planck equation, the standard method is to invoke the separation *ansatz*,

$$P(\mathbf{x}, t) = \sum_{k=0}^{\infty} \varphi_k(\mathbf{x}) e^{-\lambda_k t} \quad (8.10)$$

leading to the eigenvalue problem for the operator

$$\mathcal{L}(\mathbf{x})\varphi_k(\mathbf{x}) = -\lambda_k \varphi_k(\mathbf{x}), \quad (8.11)$$

and the associated eigenvalue problem for its adjoint

$$\mathcal{L}^\dagger(\mathbf{x})\varphi_k^*(\mathbf{x}) = -\lambda_k^* \varphi_k^*(\mathbf{x}). \quad (8.12)$$

Ordering the eigenvalues by their real parts, $\text{Re}(\lambda_0) \leq \text{Re}(\lambda_1) \leq \dots$, the existence of a steady-state to Eq. (8.10) clearly requires $\lambda_0 = 0$, with the associated eigenfunction corresponding to the steady-state distribution $\varphi_0(\mathbf{x}) = P(\mathbf{x})$. Note that when detailed balance is satisfied and the operator is self-adjoint, all eigenvalues are purely real and non-degenerate.

As mentioned in Chap. 2 in the equilibrium case and as can be seen from Eq. (8.10), the probability distribution relaxes exponentially to this steady-state at long times, with a characteristic time

$$\tau := \frac{1}{\text{Re}(\lambda_1)}. \quad (8.13)$$

Importantly, the spectral gap $\text{Re}(\lambda_1) - \text{Re}(\lambda_0) = \text{Re}(\lambda_1)$ can only be increased by the presence of out-of-equilibrium currents [281], see also [282] for a more intuitive discussion in the discrete Markov jump context. As a result, $\tau \leq \tau_{\text{eq}}$, and the NESS can therefore be reached quicker than the equilibrium steady-state. We will shortly see that complementary bounds can be devised, allowing one to understand the role of these nonequilibrium currents more precisely.

In addition to this improvement on the convergence time, it has also been shown that the introduction of out-of-equilibrium currents leads to a reduction in the asymptotic variance of arbitrary observables [283, 284].

8.1.3 Barrier crossing out-of-equilibrium

When we introduced the relaxation time in Chap. 2, we immediately mentioned that it is dominated by the crossing of the highest energy barrier when the temperature is small. For reversible Langevin dynamics, the most precise statement is the Eyring-Kramers formula, which states that the average transition time between two wells with minima located at $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$ asymptotically follows

$$\langle \tau_{\bar{\mathbf{x}}_1 \rightarrow \bar{\mathbf{x}}_2}^{\text{eq}} \rangle \underset{\beta \rightarrow \infty}{\sim} \frac{2\pi}{\lambda_+^{*,\text{eq}}} \sqrt{\frac{|\det H(\bar{\mathbf{x}}_*)|}{\det H(\bar{\mathbf{x}}_1)}} \exp(\beta \Delta U), \quad (8.14)$$

where $\bar{\mathbf{x}}_*$ is the location of the saddle point in between the two minima, which has a potential difference of $\Delta U = U(\bar{\mathbf{x}}_*) - U(\bar{\mathbf{x}}_1)$ with the starting point $\bar{\mathbf{x}}_1$. $H(\mathbf{x})$ refers to the Hessian of the potential U evaluated at the point $\mathbf{x}$, and $\lambda_+^{*,\text{eq}}$ is the magnitude of the negative eigenvalue of this Hessian at the saddle point (we assume there is only one unstable direction at the saddle). Notice that the details regarding the specific wells considered are relegated to the prefactor, meaning the dominating transition time is solely determined by ΔU .

This Eyring-Kramers formula has been rigorously extended to irreversible Langevin equations [285], see also [286, 287]. In this out-of-equilibrium context, we must define the *quasipotential* V which generalizes the role of the potential function U . Formally, it is defined as the minimum action between a stable fixed point of the force field $\mathbf{x}_1$ and an arbitrary point $\tilde{\mathbf{x}}$,

$$V(\mathbf{x}_1, \tilde{\mathbf{x}}) = \inf_{\{\mathbf{x}(t)\}} S[\mathbf{x}] \quad \text{u.c.} \quad \lim_{t \rightarrow -\infty} \mathbf{x}(t) = \mathbf{x}_1, \quad \lim_{t \rightarrow \infty} \mathbf{x}(t) = \tilde{\mathbf{x}}, \quad (8.15)$$

with the Freidlin-Wentzell action functional

$$S[\mathbf{x}] = \int_{-\infty}^{\infty} dt [\dot{\mathbf{x}} - \mathbf{F}(\mathbf{x})]^2, \quad (8.16)$$

which can be derived by considering the generating functional of noisy trajectories, similarly to what we have done in Sec. 7.5.4 of the previous chapter.

In the particular case considered here where the steady-state distribution is unaltered by the out-of-equilibrium currents, the generalized Eyring-Kramers equation yields

$$\langle \tau_{\mathbf{x}_1 \rightarrow \mathbf{x}_2} \rangle \underset{\beta \rightarrow \infty}{\sim} \frac{2\pi}{\lambda_+^*} \sqrt{\frac{|\det H(\mathbf{x}_*)|}{\det H(\mathbf{x}_1)}} \exp(\beta V(\mathbf{x}_1, \mathbf{x}_*)), \quad (8.17)$$

where now $\mathbf{x}_1$ and $\mathbf{x}_2$ are stable *fixed points* of the total force $\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}) + \mathbf{F}_{\text{neq}}(\mathbf{x})$, while $\mathbf{x}_*$ is the saddle point of the quasipotential separating the basins of attraction of $\mathbf{x}_1$ and $\mathbf{x}_2$ [285]. Importantly, λ_+^* is also redefined as the magnitude of the negative eigenvalue of the *Jacobian* of $\mathbf{F}$ evaluated at $\mathbf{x}_*$.

Clearly, the topology of the modified force field will be of central importance for potential improvements of the relaxation time at low temperatures. In any case, while one must keep in mind that the generalized Eyring-Kramers formula only formally holds for vanishing temperatures, this scaling will be useful to understand the effectiveness (or lack-off) of possible out-of-equilibrium contributions at small temperatures.

8.2 Theoretical bounds on the correlation time

Before jumping into a more precise study of specific steady-state preserving out-of-equilibrium currents, it is interesting to further study the nonequilibrium steady-states of interest in full generality. In this section, we adapt and detail the derivation of the bounds presented in [4].

So far, we have defined the relaxation time as the inverse of the real part of the first non-zero eigenvalue of the Fokker-Planck operator associated to the dynamics. While this is a mathematically precise approach to the relaxation dynamics, the diagonalization of the operator is, in general, a very difficult task that cannot be performed systematically.

Instead, we now consider the (de)correlation time specific to a given observable z . We define

$$\tau^z := \frac{\int_0^\infty dt \text{Cov}(z(t), z(0))}{\text{Var}(z(t))}. \quad (8.18)$$

Intuitively, τ^z can be understood as the typical time on which the correlations of $z(t) = z(\mathbf{x}(t))$ decay, and therefore characterizes the time required for the self-averaging of an observable along a trajectory.

8.2.1 General variational statement

A common approach in stochastic thermodynamics – the study of averages over individual, possibly out-of-equilibrium, trajectories rather than ensembles – is to use rigorous variational principles to derive upper or lower bounds on various

Chapter 8. Steady-state preserving out-of-equilibrium currents

quantities. The celebrated Thermodynamic Uncertainty Relation [288,289], which limits the square of the average current from above by the variance of current times the entropy production, is an example of such bound.

The starting point to establish a variational relation for the correlation time τ^z is to introduce the time-integrated observable

$$Z_{\mathcal{T}} = \int_0^{\mathcal{T}} dt z(\mathbf{x}(t)). \quad (8.19)$$

To characterize the fluctuations of this time integral, we also introduce its rescaled cumulant generating function

$$k^Z(h) = \lim_{\mathcal{T} \rightarrow \infty} \frac{1}{\mathcal{T}} \log \langle e^{hZ_{\mathcal{T}}} \rangle. \quad (8.20)$$

Generalizing the known variational formulae for time-integrated *currents* [290,291] to this time-integrated observables at hand, we may write

$$k^Z(h) = \sup_{\mathbf{y}} \Phi^z[\mathbf{y}], \quad \text{with} \quad \Phi^z[\mathbf{y}] = h \langle z \rangle^{\mathbf{y}} - \frac{1}{4} \beta \langle \mathbf{y}^2 \rangle^{\mathbf{y}}, \quad (8.21)$$

where $\langle \cdot \rangle^{\mathbf{y}}$ refers to an average taken with respect to the steady-state distribution of the original Langevin dynamics with the added drift $\mathbf{y}$,

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}(t)) + \mathbf{y}(\mathbf{x}(t)) + \sqrt{2T} \boldsymbol{\xi}(t). \quad (8.22)$$

Now, for any given steady-state $P^y(\mathbf{x})$, we can find a unique drift vector $\mathbf{y}^*(\mathbf{x})$ that maximizes $\Phi^z[\mathbf{y}]$, with the idea of eventually maximizing with respect to this steady-state distribution. Proceeding to the maximization on $\mathbf{y}$ under the constraint that $P^y(\mathbf{x})$ is a steady-state solution to the Fokker-Planck equation associated to Eq. (8.22) and rearranging terms, one eventually finds

$$\Phi^z[\mathbf{y}^*] = h \langle z e^{\chi} \rangle - \frac{1}{4} T \langle (\nabla(\eta^* + \chi))^2 e^{\chi} \rangle, \quad (8.23)$$

where $\chi(\mathbf{x}) = \log \left(\frac{P^y(\mathbf{x})}{P(\mathbf{x})} \right)$ and $\eta^*(\mathbf{x})$ is essentially a Lagrange multiplier enforcing the constraint rewritten as

$$\nabla \cdot \left[P(\mathbf{x}) e^{\chi(\mathbf{x})} (\mathbf{F}_{\text{neq}}(\mathbf{x}) + T \nabla \eta^*(\mathbf{x})) \right] = 0, \quad (8.24)$$

which, importantly, is the Euler-Lagrange equation of the convex minimization problem

$$\inf_{\eta} \left[\langle (T(\nabla \eta)^2 + 2 \nabla \eta \cdot \mathbf{F}_{\text{neq}}) e^{\chi} \rangle \right]. \quad (8.25)$$

Further manipulating the terms in the right hand side of Eq. (8.23) and maximizing with respect to $P^y(\mathbf{x})$, or equivalently $\chi(\mathbf{x})$ provided $e^{\chi(\mathbf{x})}P(\mathbf{x})$ is a normalized probability density, the rescaled cumulant generating function can finally be rewritten as

$$k^Z(h) = \sup_{\chi} \inf_{\eta} \left[\frac{1}{\langle e^{\chi} \rangle} \left(h \langle z e^{\chi} \rangle + \langle \nabla \eta \cdot \mathbf{F}_{\text{neq}} e^{\chi} \rangle + T \langle (\nabla \eta)^2 e^{\chi} \rangle - \frac{1}{4} T \langle (\nabla \chi)^2 e^{\chi} \rangle \right) \right]. \quad (8.26)$$

We can now use this variational expression to find bounds on functions of the cumulants of $Z_{\mathcal{T}}$. Taylor expanding the rescaled cumulant generating function, we indeed have

$$k^Z(h) \approx h \langle z \rangle + \frac{1}{2} h^2 \mathcal{T} \text{Var}(\bar{z}_{\mathcal{T}}) + O(h^3), \quad (8.27)$$

where we have identified the time average

$$\bar{z}_{\mathcal{T}} = \frac{1}{\mathcal{T}} \int_0^{\mathcal{T}} dt z(\mathbf{x}(t)), \quad (8.28)$$

and assumed that the system is ergodic i.e.

$$\lim_{\mathcal{T} \rightarrow \infty} \bar{z}_{\mathcal{T}} = \langle z \rangle. \quad (8.29)$$

Rescaling $\chi \rightarrow h\chi$, $\eta \rightarrow h\eta$ and similarly expanding Eq. (8.26) in powers of h , we find that the linear contributions in h cancel out while equating second order terms yields

$$\frac{1}{2} \mathcal{T} \text{Var}(\bar{z}_{\mathcal{T}}) = \sup_{\chi} \inf_{\eta} \left[\text{Cov}(z, \chi) + \langle \chi \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle + T \langle (\nabla \eta)^2 \rangle - \frac{1}{4} T \langle (\nabla \chi)^2 \rangle \right]. \quad (8.30)$$

Finally, going back to the original interpretation of τ^z , $z(\mathbf{x}(t))$ can be considered to be independent to $z(\mathbf{x}(t + \mathcal{T}))$ for $\mathcal{T} > \tau^z$. As a result, the time average can be understood as a sum of independent random variables, in which case by a central limit theorem argument (subdividing $\mathcal{T}$ in independent chunks of length $2\tau_z$),

$$\frac{\tau^z}{\mathcal{T}} \approx \frac{1}{2} \frac{\text{Var}(\bar{z}_{\mathcal{T}})}{\text{Var}(z)} \quad (8.31)$$

for $\mathcal{T} \gg \tau^z$. Linearly rescaling χ and η and subsequently solving the extremizing problem on their respective factors, the final variational formula for τ^z reads

$$\tau^z = \frac{1}{\text{Var}(z)} \sup_{\chi} \left[\frac{\text{Cov}(z, \chi)^2}{T \langle (\nabla \chi)^2 \rangle + \sup_{\eta} f(\eta, \chi)} \right], \quad f(\eta, \chi) = \frac{\langle \chi \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle^2}{T \langle (\nabla \eta)^2 \rangle}. \quad (8.32)$$

This is a rather technical statement, and the previously summarized derivation is quite involved. Nonetheless, we will see that it provides an essential basis to

Chapter 8. Steady-state preserving out-of-equilibrium currents

derive more intuitive bounds that significantly clarify the role of the different properties of detailed balance violating currents.

Before doing so, let us discuss some points relative to this bound in the context of equilibrium dynamics. First, in the spirit of the result on the spectral gap of the operator, notice that $\tau^z \leq \tau_{\text{neq}}^z$, as the out-of-equilibrium contribution in the denominator can only decrease the object that we aim to minimize. Note also that the variational expression allows us to derive a general bound in the equilibrium case $\mathbf{F}_{\text{neq}}(\mathbf{x}) = \mathbf{0}$ by taking $\chi(\mathbf{x}) = z(\mathbf{x})$, resulting in

$$\tau_{\text{eq}}^z \geq \frac{\text{Var}(z)}{T\langle(\nabla z)^2\rangle}. \quad (8.33)$$

While seemingly purely mathematical, this upper limit on the decorrelation time can be understood rather intuitively. Indeed, it can be shown that the short-time fluctuations of the displacement are given by

$$\lim_{\delta t \rightarrow 0} \frac{1}{2} \frac{\text{Var}(\delta z)}{\delta t} = T\langle(\nabla z)^2\rangle, \quad \text{where} \quad \delta z = z(\mathbf{x}(t + \delta t)) - z(\mathbf{x}(t)). \quad (8.34)$$

As a result, the bound of Eq. (8.33) implies that the correlation time is limited by a tradeoff between short and long time fluctuations of the observable. Finally, consider the *intrinsic* correlation time defined as

$$\tau^* := \sup_z [\tau^z], \quad (8.35)$$

which can be understood as the time required for the slowest observable to decorrelate. First taking the supremum on z , we find

$$\tau_{\text{eq}}^* = \sup_{\chi} \left[\frac{\text{Var}(\chi)}{T\langle(\nabla \chi)^2\rangle} \right]. \quad (8.36)$$

Remarkably, the right hand side is a known variational formula for the smallest non-zero eigenvalue of the Fokker-Planck operator when detailed balance is satisfied [92], meaning we recover the relaxation time,

$$\tau_{\text{eq}}^* = \frac{1}{\lambda_1}. \quad (8.37)$$

This equality thus formally establishes a link between the relaxation of the probability distribution and the time required for the slowest observable in the system to decorrelate. Unfortunately, no such variational expression exists for the eigenvalues of the Fokker-Planck operator in the irreversible case.

8.2.2 Dissipation speed limit

When the system is not in detailed balance, it is desirable to relate the improvement in the relaxation to the rate of entropy production, which quantifies the violation of time-reversal symmetry.

Going back to the variational formula for τ^z , we may notice that

$$\langle \chi \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle^2 = \langle (\chi - \chi_0) \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle^2 \quad (8.38)$$

when χ_0 is a constant, as $\langle \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle = 0$ from integration by parts and the global balance condition. By the Cauchy-Schwarz inequality, we then have

$$\langle \chi \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle^2 \leq \frac{1}{T} \langle (\chi - \chi_0)^2 \mathbf{F}_{\text{neq}}^2 \rangle T \langle (\nabla \eta)^2 \rangle, \quad (8.39)$$

and as a result,

$$f(\eta, \chi) \leq \frac{1}{T} \langle (\chi - \chi_0)^2 \mathbf{F}_{\text{neq}}^2 \rangle = \langle \chi^2 \tilde{\sigma} \rangle - 2\chi_0 \langle \chi \tilde{\sigma} \rangle + \chi_0^2 \sigma, \quad (8.40)$$

where we have defined the *local* rate of entropy production $\tilde{\sigma}(\mathbf{x}) \equiv \frac{1}{T} \mathbf{F}_{\text{neq}}^2(\mathbf{x})$ such that $\sigma \equiv \langle \tilde{\sigma} \rangle$. Maximizing over the constant χ_0 , the equality of Eq. (8.32) becomes an inequality,

$$\tau^z \geq \frac{1}{\text{Var}(z)} \sup_{\chi} \left[\frac{\text{Cov}(z, \chi)^2}{T \langle (\nabla \chi)^2 \rangle + \langle \chi^2 \tilde{\sigma} \rangle - \frac{1}{\sigma} \langle \chi \tilde{\sigma} \rangle^2} \right]. \quad (8.41)$$

This bound can finally be loosened to improve its interpretation by arbitrarily taking $\chi(\mathbf{x}) = z(\mathbf{x})$, resulting in

$$\tau^z \geq \frac{\text{Var}(z)}{T \langle (\nabla z)^2 \rangle + \sigma \text{Var}_{\sigma}(z)}, \quad (8.42)$$

where $\text{Var}_{\sigma}(\cdot)$ refers to the variance taken with respect to the entropy-rescaled probability density $P_{\sigma}(\mathbf{x}) := \frac{\tilde{\sigma}(\mathbf{x})}{\sigma} P(\mathbf{x})$.

We refer to this bound as the dissipation speed limit. Indeed, it demonstrates that while the dynamics of the system are accelerated by driving it out-of-equilibrium, there is a minimum amount of dissipation, characterized by the entropy production rate, associated to this acceleration.

A very interesting byproduct of this speed limit is that it may be used to provide a lower bound on the rate of entropy production by measuring the correlation time, see [4]. The novelty of such an approach is that the dissipation may be estimated even from a time-symmetric observables when there are no observable real-space currents. This possibility is particularly interesting in models, e.g. of

active matter, where the TRS violation may be apparent only in hidden degrees of freedom, preventing its direct observation, see [292] and references therein.

Note that a very similar bound, which may also be used to estimate the rate of entropy production, can be established in the frequency domain [293]. Whereas a reversible process displays a monotonously decaying power spectral density (PSD) for the fluctuations of any observable z , the introduction of nonequilibrium currents generically leads to the emergence of a peak at a finite frequency. It is then the magnitude of this peak relative to the high frequency asymptote of the PSD that can be bounded and related to the rate of entropy production.

8.2.3 Geometric speed limit

Alternatively, following an integration by parts, we may notice that

$$\langle \chi \nabla \eta \cdot \mathbf{F}_{\text{neq}} \rangle^2 = \langle \eta \nabla \chi \cdot \mathbf{F}_{\text{neq}} \rangle^2. \quad (8.43)$$

Restricting ourselves to fields $\chi(\mathbf{x}) = \chi_{\perp}(\mathbf{x})$ the gradient of which are orthogonal to the out-of-equilibrium currents, i.e. satisfying $\nabla \chi_{\perp}(\mathbf{x}) \cdot \mathbf{F}_{\text{neq}}(\mathbf{x}) = 0$, the above equality allows us to write the alternative bound

$$\tau^z \geq \frac{1}{\text{Var}(z)} \sup_{\chi_{\perp}} \left[\frac{\text{Cov}(z, \chi_{\perp})^2}{T \langle (\nabla \chi_{\perp})^2 \rangle} \right]. \quad (8.44)$$

This geometric speed limit is now importantly independent of the specific detailed balance violation and its strength. As a result, we conclude that one cannot accelerate the dynamics arbitrarily much simply by increasing the dissipation and therefore the rate of entropy production in the system.

Intuitively, such a result can be interpreted as follows. Going back to the correspondance between the largest correlation time and the eigenvalue gap of the Fokker-Planck operator, the relaxation in an equilibrium system is limited by its slowest timescale. Driving the system out-of-equilibrium first allows one to speedup this bottleneck process, at the cost of introducing dissipation. At some point, however, another process that originally was associated to a faster timescale and is unaffected by the driving³⁶ becomes the bottleneck. This new bottleneck being immune to the effects of the out-of-equilibrium currents, further increasing the driving strength and dissipation will not improve the situation, unless the geometric structure of the driving is also altered.

³⁶If the observable itself satisfies $\nabla z(\mathbf{x}) \cdot \mathbf{F}_{\text{neq}}(\mathbf{x}) = 0$, then the geometric speed limit is $\tau^{z_{\perp}} \geq \frac{\text{Var}(z_{\perp})}{T \langle (\nabla z_{\perp})^2 \rangle}$, corresponding to the equilibrium bound from Eq. (8.33), meaning there is no speedup for such observables.

8.3 Accelerating relaxation in practice

We have determined that having a non-zero detailed balance violating component $\mathbf{F}_{\text{neq}}(\mathbf{x})$ may be beneficial to speedup convergence. At this stage, however, it is not clear what specific choice to take given the very general form of the criterion given in Eq. (8.5). Building on existing methods and the insights provided by the speed limits derived above, we will now attempt to understand why specific force fields may or may not be effective to accelerate the dynamics.

8.3.1 Ichiki-Ohzeki forcing

Let us start with a simple yet important observation. Placing ourselves in a single dimension, the constraint on the nonequilibrium force becomes

$$\frac{d}{dx}(F_{\text{neq}}(x)P(x)) = 0, \quad (8.45)$$

leaving $F_{\text{neq}}(x) \propto e^{\beta U(x)}$ as the only possibility. Not only is this somewhat restrictive, but it is also rather prohibitive from the numerical standpoint. Indeed, having a force of exponential magnitude will likely require a very fine discretization in time, meaning that the improvement provided by the driving might be overshadowed by a possibly costly time step requirement. As a matter of fact, the same holds for the trivial forcing $\mathbf{F}_{\text{neq}}(\mathbf{x}) \propto \mathbf{1}e^{\beta U(\mathbf{x})}$ in higher dimensions.

As an alternative for all dimensions (including $d = 1$), Ichiki & Ohzeki (IO) proposed to take two copies of an identical equilibrium system and to introduce a detailed balance violating coupling between the two [294]. More precisely, the dynamics follow

$$\begin{aligned} \dot{\mathbf{x}}^{(1)}(t) &= -\nabla_{\mathbf{x}^{(1)}}U(\mathbf{x}^{(1)}(t)) + \gamma\nabla_{\mathbf{x}^{(2)}}U(\mathbf{x}^{(2)}(t)) + \sqrt{2T}\boldsymbol{\xi}^{(1)}(t) \\ \dot{\mathbf{x}}^{(2)}(t) &= -\nabla_{\mathbf{x}^{(2)}}U(\mathbf{x}^{(2)}(t)) - \gamma\nabla_{\mathbf{x}^{(1)}}U(\mathbf{x}^{(1)}(t)) + \sqrt{2T}\boldsymbol{\xi}^{(2)}(t), \end{aligned} \quad (8.46)$$

where the noise terms are taken to be entirely independent and γ is a coupling constant setting the strength of the nonequilibrium currents.

The change of sign in front of the second terms on the right hand sides leads the combined force vector to be divergence-free with respect to the combined degrees of freedom, while its orthogonality with the gradient of the potential ensures that Eq.(8.6) is satisfied. In more intuitive terms, this means that the out-of-equilibrium currents are purely non-gradient and *along equiprobability lines* of the joint probability distribution of the two copies. Such currents are represented in a simplified setting in Fig. 8.1. As a result, one can show that (i) the fixed points of the modified force field are identical to that of the equilibrium dynamics, and

therefore remain the extrema of the potential and (ii) that the quasipotential between a minimum and its associated saddle point is equal to the potential difference of the duplicated equilibrium problem [295]. In addition to these properties, the very constrained geometry of the force suggests that the geometrical bound of Eq. (8.44) is quickly reached, as many observables will likely not benefit from the acceleration that is essentially limited to “stirring” inside each well.³⁷

With these specificities in mind, we can go back to the generalized Eyring-Kramers relation given in Eq. (8.17). Due to points (i) and (ii) raised above, the only difference between the equilibrium and out-of-equilibrium escape times lies in λ_+^* . As pointed out in Ref. [295], assuming the trajectory of interest leaves a *stiff* well (i.e. for which the determinant of the Hessian of U is large), then in the small driving limit the scaling suggests

$$\tau \sim \frac{1}{1 + \gamma^2} \tau_{\text{eq}}. \quad (8.47)$$

How effective are these IO currents to accelerate relaxation in practice? In their original publication, Ichiki & Ohzeki reported a very significant acceleration in the sampling of a simple one-dimensional double well and, importantly, an elimination of the critical slowing down of the XY model for $\gamma = 10$ [294]. To accommodate for this large forcing, however, the stochastic dynamics must be implemented using a Heun scheme, as the standard explicit Euler-Murayama method is insufficiently precise.

Perhaps more interestingly in the context of this thesis largely dedicated to disordered systems, Ghimenti & Van Wijland proposed an implementation of the scheme in the mean-field ($N \rightarrow \infty$) dynamics of the spherical p -spin [295], which is known to display ergodicity breaking below the so-called dynamical temperature T_d for $p \geq 3$.³⁸ Above T_d , the authors find an appreciable reduction in the relaxation time, with an improvement scaling as Eq. (8.47) for $\gamma \lesssim O(1)$ (but expected to saturate at some point). At lower temperatures, however, the forcing is unfortunately insufficient to shift the value of T_d and the dynamics remain frozen below this unchanged dynamical temperature in the $p = 3$ case tested.

This being said, one must keep in mind that the spherical p -spin for $p \geq 3$ is a very harsh test for the Ichiki-Ohzeki method. Indeed, at the dynamical temperature T_d , the free energy landscape of the model goes from having a single minimum to an exponential number of metastable states, i.e. a finite complexity. Moreover, these minima are separated by barriers that scale with N and therefore diverge in the $N \rightarrow \infty$ limit, leading to true ergodicity breaking. As such, and as we have seen that divergence-free currents only affect one of the prefactors of the

³⁷Note that this was already pointed out in Ref. [295].

³⁸The spherical p -spin is the generalization to order p interactions of the spherical SK model mentioned in Chap. 2, which corresponds to $p = 2$.




Figure 8.1: Total force fields $\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}) + \mathbf{F}_{\text{neq}}(\mathbf{x})$ for the $2d$ double well potential of Eq. (8.51), colors given by the value of the potential (lighter corresponding to smaller) for (a) original equilibrium system; (b) standard Ichiki-Ohzeki forcing with $\gamma = 5$. Light stars, dark stars and dark squares correspond to stable fixed points, saddle points (single unstable direction) and unstable fixed points of the force field respectively.

generalized Eyring-Kramers law, it is somewhat natural that Ichiki-Ohzeki forcing is unable to overcome such energy barriers in the mean-field dynamics.

Note that generically this type of divergence-free current can be constructed without having to duplicate the system when $d \geq 2$. Indeed, introducing a $d \times d$ skew-symmetric matrix $\mathbf{J}^\top = -\mathbf{J}$, one can take

$$\mathbf{F}_{\text{neq}}(\mathbf{x}) = \gamma \mathbf{J} \nabla U(\mathbf{x}), \quad (8.48)$$

which is also divergence-free and orthogonal to the probability density gradient. In contrast with the standard Ichiki-Ohzeki prescription, the choice of the matrix $\mathbf{J}$ is not very constrained, and picking a particular structure can therefore be awkward in high dimensions.

8.3.2 An extension to complicated currents

To go beyond divergence-free currents, a recent proposition formulated by Dechant [296] is to take the *temperature dependent* out-of-equilibrium force term

$$\mathbf{F}_{\text{neq}}^T(\mathbf{x}) = \gamma(\mathbf{x}) \mathbf{J} \nabla U(\mathbf{x}) - T \mathbf{J} \nabla \gamma(\mathbf{x}), \quad (8.49)$$

where $\gamma(\mathbf{x})$ is now a detailed balance violating coupling *field*. For a constant $\gamma(\mathbf{x})$, one immediately recovers the divergence-free choice mentioned above, while this type of forcing can also be understood in an Ichiki-Ohzeki way by taking $\mathbf{x} = [\mathbf{x}^{(1)} \quad \mathbf{x}^{(2)}]^\top$ and subsequently adapting the gradient terms and the skew symmetric matrix.

Chapter 8. Steady-state preserving out-of-equilibrium currents

In the $\beta \rightarrow \infty$ limit, the contribution that is not orthogonal to the gradient of the potential vanishes, as required by Eq. (8.6). As a result, the effect of the current is expected to be identical to the standard Ichiki-Ohzeki method at very small temperatures. Note that this is consistent with intuition: without thermal fluctuations, having currents going across iso-probability lines would undoubtedly change the steady-state distribution. At small but finite temperatures, however, the fixed points of this new force field do not necessarily coincide with the minima and saddles of the potential. As a result, we expect that the entire generalized Eyring-Kramers equation can be modified when the temperature is small but still finite, and not only the eigenvalue of the Jacobian as was the case above. At first glance, the fact that the current is no longer restricted to be along equiprobability lines appears promising: very pictorially, the out-of-equilibrium currents could perhaps help the system “hike” its way out of a well!

In the following, we will focus on the $d = 2$ case, which is easy to visualize and understand in terms of gradient and rotational contributions, and take

$$\mathbf{J} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad (8.50)$$

to be as close as possible to the Ichiki-Ohzeki prescription. Considering a two-dimensional double-well potential,

$$U(\mathbf{x}) = \sum_{i=1}^2 \left(\frac{1}{4} x_i^4 - \frac{1}{2} x_i^2 \right), \quad (8.51)$$

we have illustrated the standard IO forcing in Fig. 8.1 and more complicated temperature-dependent currents in Fig. 8.2.

A first interesting option offered by this extension to a coupling field is to only enforce a strong driving where it actually has a significant effect on the relaxation dynamics. As highlighted by the previously derived geometric bound, we indeed know that brutally driving the system will eventually not improve the situation in general. More specifically, the Eyring-Kramers equation has allowed us to identify that the standard Ichiki-Ohzeki forcing only has a sizable effect due to its influence at the saddle point. In Fig. 8.2(a), we therefore consider the forcing field

$$\gamma(\mathbf{x}) = \gamma \sum_{\{\mathbf{x}_*\}} e^{-\frac{(U(\mathbf{x}) - U(\mathbf{x}_*))^2}{U_0}}, \quad (8.52)$$

where the sum is over all the saddle points of the potential and U_0 is an adjustable parameter parametrizing the confinement of the forcing to the saddle points. Of course, such a prescription requires the knowledge of the saddle points, however one could imagine a heuristic weight based on the numerical exploration of the




Figure 8.2: Total force fields $\mathbf{F}(\mathbf{x}) = -\nabla U(\mathbf{x}) + \mathbf{F}_{\text{neq}}(\mathbf{x})$ for the $2d$ double well potential of Eq. (8.51), colors given by the value of the potential (lighter corresponding to smaller) for generalized Ichiki-Ohzeki forcing with (a) the targeted coupling field given in Eq. (8.52), $\gamma = 5$, $T = 0.1$; (b) an arbitrary space dependent coupling field $\gamma(\mathbf{x}) = \gamma \sum_{\mathbf{x}_*} e^{-\frac{(\mathbf{x}-\mathbf{x}_*)^2}{\kappa}}$, $\gamma = 5$, $\kappa = T = 0.1$. Light stars, dark stars and dark squares correspond to stable fixed points, saddle points (single unstable direction) and unstable fixed points of the force field respectively.

model or analytical calculations. The main advantage here would be to reduce the computational constraint on the step size that can result from large forcing. Even conceptually, this is an interesting example to illustrate that a simple increase in the rate of entropy production (resulting from a homogeneous forcing in the entire system) will not improve the effects of the driving due to geometrical limitations. Numerical experiments confirm the equivalent effectiveness of this targeted IO current relative to the original prescription (Fig. 8.3(b)).

To obtain more interesting and complicated currents, one needs to decouple the forcing field from the potential. As mentioned above, doing so at sufficiently large temperatures, we can create force fields that have new or modified fixed points. Remarkably, this also includes stable fixed points, as illustrated in Fig. 8.2(b). As a result of this much stronger modification, the transition path between two minima no longer has to go through the saddle point of the original potential. This phenomenon is illustrated with numerically simulated trajectories in Fig. 8.3(a). This being said, engineering an effective force field is not an easy task. Besides, what the thermodynamic bounds on the correlation time has demonstrated is that the acceleration will not affect all observables equally. At this time, we have not managed to find a non-trivial force field that clearly does better than the standard Ichiki-Ohzeki method in terms of the numerically observed relaxation time, see Fig. 8.3(b). At low but finite temperatures, it nonetheless appears unlikely that the IO driving field is optimal – we have in fact shown that it is not for a given amount of entropy production.




Figure 8.3: (a) Sample of trajectories going from one well to another in almost directly (here leaving $\mathbf{x}(0) = [-1 \ 1]^T$ and reaching a small circle around $\mathbf{x} = [-1 \ -1]^T$) for the complicated force field illustrated in Fig. 8.2(b), albeit for $T = 0.3$. The black star represents the saddle point of the original potential, which is clearly not visited by the most probable trajectories in the modified force field. (b) Comparison of the reduction in relaxation time (measured by fitting the relaxation dynamics of the position) for the original IO forcing and the two force fields of Fig. 8.2 ($T = 0.1$). The continuous line represents the full Eyring-Kramers prediction of the standard IO acceleration.

8.3.3 Perspectives

In order to improve our understanding of what constitutes a good topology for accelerating relaxation, studying the most probable transition path and the associated quasipotential at low but finite temperature appears to be the natural next step. This being said, computing this trajectory when the temperature is finite and currents much richer is not entirely straightforward. Indeed, the standard *instanton*, i.e. the most likely trajectory between two points in the vanishing temperature limit, is found by minimizing the previously introduced Freidlin-Wentzell action of Eq. (8.16) for $T = 0$ strictly. At finite temperatures, on the other hand, we should also consider the effects of thermal fluctuations.

Formally, one must study the probability to find a trajectory within a small tube, as the probability of a single trajectory at finite temperature is always zero [297, 298]. The most likely infinitesimal tube of trajectories can then be shown to be that minimizing the Onsager-Machlup action

$$S^T[\mathbf{x}] = \int_{-\infty}^{\infty} dt ([\dot{\mathbf{x}} - \mathbf{F}(\mathbf{x})]^2 + T \nabla \cdot \mathbf{F}(\mathbf{x})). \quad (8.53)$$

The new temperature-dependent term in the action has been shown to be essential to recover the correct transition paths, including with laboratory experiments [299]. As a matter of fact, this term is likely of even greater importance for

our generalized IO currents, as it is of the same order in temperature as the leading contribution stemming from the $\nabla\gamma$ term in Eq. (8.49). Beyond the strict numerical minimization of S^T (e.g. using the method detailed in [275, 300–302] or [303]), a possible and interesting extension is to study the distribution of these paths. At finite temperature, the transition paths are indeed formally drawn from a Gibbs-Boltzmann distribution $P[\{\mathbf{x}(t)\}] \propto \exp(-\beta S^T[\mathbf{x}])$. One can then for instance sample these paths by following the evolution of trajectories in a fictitious time with a functional Langevin equation, where $S_T[\mathbf{x}]$ now plays the role of the potential – an approach coined *metadynamics* in [304].

In any case, once the feasibility of effective complicated currents is determined (with minimum action paths or otherwise), it would be extremely interesting to assess if these non-trivial force fields could now alter the low temperature ergodicity breaking phenomena displayed by glassy models. In mean-field spin-glasses such as the spherical p -spin discussed above, the divergence of energy barriers suggests that the freezing in the low temperature phase is unavoidable, but simpler trap models may be somewhat less challenging.

In two dimensions, where the physical understanding of the driving appears more accessible than in high-dimensional systems, an appropriate test case for the effectiveness of a given force field (or a general construction method) could be that of a Gaussian random potential $\phi(\mathbf{x})$. Such a random potential is usually taken to be of zero mean, and is characterized by its correlation function,

$$\overline{\phi(\mathbf{x})\phi(\mathbf{x}')} = \Gamma(|\mathbf{x} - \mathbf{x}'|). \quad (8.54)$$

Gaussian potentials are indeed known to display “super Arrhenius” behavior [305], meaning the relaxation time averaged over the disorder scales as

$$\overline{\tau^{\text{eq}}} \sim \exp(\beta^2 A). \quad (8.55)$$

Beating the Arrhenius scaling at low but finite temperature using an appropriately designed current would therefore provide a very sizable improvement to the diffusion at large scales in this type of random potential. A more extreme test would be to take the squared Gaussian potential $\psi(\mathbf{x}) = \frac{1}{2}\phi(\mathbf{x})^2$. Unlike the standard Gaussian in which the dynamics are only frozen at strictly zero temperature, the squared Gaussian potential is a *bona fide* trap model with a dynamical transition temperature $T_d = 1$ below which $\overline{\tau^{\text{eq}}}$ diverges [306].

8.4 Conclusion

In this chapter, we have taken a rather significant sidestep from the socioeconomic leitmotiv of the thesis. Indeed, all the problems we had studied so far were characterized by strong interactions: between assets in the portfolio problem, between

individual agents in the interacting version of the Slutsky problem, the SK-game and the Sakoda-Schelling model. Here, we have considered a much more simple setup, with a single particle following the standard overdamped Langevin equation. Yet, doing so has importantly allowed us to isolate the impact of detailed balance violating contributions to the dynamics. Enforcing that the steady-state distribution remains unchanged indeed provides a somewhat unique opportunity to compare the equilibrium and out-of-equilibrium relaxations on an equal footing. Given the common occurrence of detailed balance violation in agent-based models as a direct consequence of individualistic decision-making rules, improving our understanding of out-of-equilibrium driving in such a simple case can only help our intuition regarding more complex systems.

It has long been known that the introduction of irreversible currents generically reduces the relaxation time, offering a straightforward application to steady-state preserving out-of-equilibrium currents in numerical sampling. With this potential use in mind, we have derived more precise bounds on the *correlation* time of arbitrary observables. In equilibrium, the largest correlation time of the system can formally be shown to coincide with the relaxation time to the steady-state. Although this direct link has no equivalent out-of-equilibrium, the influence of irreversible currents on the correlation time (which can also only be improved by the detailed balance violation) remains instructive.

Starting from a general variational formulation of the correlation time of a given observable, two relevant speed limits can be identified. First, a *dissipation* bound, which formally establishes that a minimal amount of dissipation, i.e. of entropy production, is required to accelerate the dynamics of the system. Second, and perhaps more counterintuitively, a *geometrical* bound, demonstrating that a given amount of dissipation will not accelerate *all* observables in the system, as some will necessarily be orthogonal to the current. As such, blindly increasing the amount of dissipation in the system will not improve the relaxation time arbitrarily much. There is necessarily a saturation at some point, caused by the fact that the slowest observable in the system eventually becomes one that is not affected by the driving.


These bounds may then be of use to support the design of out-of-equilibrium currents aimed at reducing the relaxation time towards a given Gibbs-Boltzmann distribution. The geometric bound, in particular, suggests that the standard Ichiki-Ohzeki method, which relies on the introduction of purely rotational currents that only act along equiprobability lines, might not be the best possible approach. At finite temperatures, it is indeed possible to extend this method to more complicated currents, for which the most probable path from an energy well to another may be strongly modified (Fig. 8.2(a)). While we have only presented very preliminary results and ideas on this problem, future research using the tools developed for vanishing temperature instantons appears to be a promising avenue.

Going back to socioeconomic systems, this research direction may turn out to be relevant on two fronts. First, from a purely computational point of view. In disordered systems described by a Gibbs-Boltzmann distribution, for example the SK game with symmetric interactions ($\varepsilon = 0$), it is indeed always interesting to accelerate relaxation in numerical experiments.³⁹ In addition to this practical consideration, a central question in the application of statistical mechanics to socioeconomic system is the robustness of equilibrium results to potentially more realistic agent-based models where detailed balance is violated. This issue, which was at the heart of Chap. 7 and discussed in Chap. 4, may indeed benefit from a precise assessment of the possible consequences of detailed balance violation on the phenomenology of relatively complex systems. Determining if the currents studied in this chapter may for instance kill the “super Arrhenius” behavior of particles diffusion in a random Gaussian potential would be a first step in this direction.

³⁹Similarly to the p -spin discussed above, we expect such acceleration to be limited in *radically complex* models such as the SK-game.

Key takeaways

- By enforcing *global balance* instead of the usual *detailed balance*, one can construct out-of-equilibrium dynamics with the same steady-state distribution as a prescribed equilibrium system.
- The presence of irreversible currents is known to *accelerate* relaxation.
- A variational expression on the *correlation time* of any arbitrary state observable allows the derivation of two distinct lower bounds:
 1. A *dissipation* bound establishes that a minimal amount of entropy production is required to achieve a given reduction of the correlation time.
 2. A *geometric* bound demonstrates that one can however not reduce the correlation time arbitrarily much simply by increasing the dissipation.
- These speed limits can be used to support the understanding of the effect of given irreversible currents, for example the saturation of the divergence-free Ichiki-Ohzeki acceleration.
- At finite temperature, the introduction of more complicated currents could lead to improvements, although what constitutes an effective geometry is still unclear at this stage.



Part V

Conclusion and future works

Conclusion

*The best that most of us can hope to achieve in physics
is simply to misunderstand at a deeper level.*

Wolfgang Pauli

Overview of the results

Let us briefly summarize what we have achieved in this thesis.

Through the study of a constrained portfolio optimization problem, we have first illustrated how simple problems may lead an agent to face radically complex situations [1]. Due to a very large number of quasi-equivalent solutions, the outcome of the risk minimization that is expected from the agent becomes extremely sensitive to the problem parameters. As a result, both rational decision making and common information can no longer be assumed.

With this in mind, we have revisited a classical problem in consumer choice theory with a boundedly rational decision rule [2]. Thanks to our specific choice of decision rule, we were able to use the powerful methods of statistical mechanics to show that bounded rationality does not necessarily lead to a breakdown of the proverbial symmetry of the Slutsky matrix, which quantifies the change in demand of a good following the change of price of another. When agents are interacting, however, we found this symmetry to be violated, in particular at the transition of a collective concentration phenomenon. Importantly, the illustrative interacting socioeconomic system that we considered also revealed an inherent limitation of equilibrium statistical mechanics to describe agent-based models. Thermodynamic equilibrium indeed requires a global free energy to be minimized by the dynamics of the system, which generally cannot be the case when agents are somewhat individualistic in their decision making.

Following these results and observations, we introduced a unifying binary choice model, the “SK-game”, showcasing bounded rationality, non-reciprocity, and learning [3]. At odds with the belief that agents learn the best response to a

situation based on past experience, we found that non-stationarity is self-induced by the complexity of the game that the agents are trying to learn. Even when the interactions are sufficiently reciprocal for the learning to converge, the outcome is far from optimal, and the agents end up in satisficing solutions. Despite its relatively simple statement, the model displays a wide range of out-of-equilibrium phenomena that have proved very challenging to understand analytically.

Departing from this high-dimensional, spin-glass inspired formalism, we considered a generalized Sakoda-Schelling occupation model [5]. As a result of the individualisation of the density perceived by the agents, the system is driven out-of-equilibrium for any non-trivial choice of the utility function and regardless of the decision making process, reiterating that detailed balance satisfying dynamics in agent-based models are the exception rather than the rule. Fortunately, the mean-field description of the model allowed us to draw parallels with the theory of active matter, which has generated much progress in the understanding of nonequilibrium dynamics in recent years. Importantly, we found the condensation of agents in suboptimal tightly packed clusters to be robust to this out-of-equilibrium setting and to decision rules beyond the usual logit choice.

Finally, taking a minor sidestep to improve our understanding of the impact of out-of-equilibrium currents on the relaxation towards a steady-state, we studied the influence of driving when it preserves the Gibbs-Boltzmann steady-state distribution. We first derived some rigorous bounds on the correlation time of arbitrary observables [4]. Using a very simple example, we then illustrated how these bounds can be used in the design of complicated out-of-equilibrium currents aiming at accelerating the relaxation of the system. Beyond the direct application of these currents to improve the numerical sampling of socioeconomic systems described by an equilibrium distribution, the way detailed balance violation may or may not impact the phenomenology of relatively complex models could provide clues regarding the robustness of some known results to decision rules beyond the standard logit choice in a more general context.

In even fewer words, the main messages of this thesis can be condensed as:

- i. The heterogeneous nature of the interactions in socioeconomic systems very likely leads to radical complexity and a highly non-trivial solution space, at odds with the common unique “equilibrium” assumption;
- ii. While equilibrium statistical mechanics may be suitable to characterize *some* socioeconomic systems, the intrinsically individualistic decision rules describing agents generically lead to out-of-equilibrium dynamics, which must then be analyzed accordingly (although the phenomenology may be relatively unchanged);
- iii. As a consequence of both points (i) and (ii), simple learning rules are unlikely

to salvage any form of rational expectations. When fixed points exist they are non-unique and lead to subtle long-time dynamics, when they do not then the agents' strategies evolve forever.

Extensions and closing remarks

Throughout the manuscript, we have mentioned some of the very numerous extensions that one could imagine for each of the problems considered. Without going into too much detail, let us mention what we believe are the most evident future directions to explore.

First and foremost, our analytical understanding of the SK-game remains limited, and therefore requires further investigations. We can notably mention the long-time behavior of the small asymmetry, low temperature region in the limit of long-term memory, as it is unclear at what stage the intentions of agents partially or entirely decorrelate in time. This question is, in our opinion, particularly important, as we have highlighted its relevance for other complex systems, biological neural networks in particular. In a similar vein, we were unable to clearly understand the influence of learning on the period of the limit cycles that emerge with perfectly rational agents in the close to zero sum regime. Recall that this region of the parameter space is particularly important conceptually, as agents spontaneously coordinate to do better on average than the Nash equilibrium, at the cost of their own instantaneous individual rewards.

Other notable immediate extensions to the present work include the computation of the quenched complexity in the portfolio problem, and the inclusion of a coupled price field in our neighborhood-less Schelling model.

In a sense, any progress towards the *clarification* and *simplification* of the fundamental understanding of these complex systems would be beneficial, in particular if we hope for these ideas to one day permeate towards mainstream economics. Indeed, the problems considered here have proved to be extremely rich in terms of statistical physics, and as a consequence many of the results are rather technical and subtle. Of course, this is what makes these interdisciplinary applications exciting: by considering new constraints and rules, we unveil new physics and make a modest contribution to the wealth of phenomena described by statistical mechanics. In doing so however, it is sometimes easy to lose sight of the original objectives and motivations. While there undoubtedly remains some aspects to expand upon and to understand more clearly, an important future effort should perhaps be to carefully communicate some of the key ideas presented here to a wider audience, in particular towards the economics community. As noted in the introduction, mainstream economics tends to be quite insular and reluctant to

new ideas, to say the least, making this endeavor sometimes difficult.⁴⁰ The very idea of toy models, which have been at the center of our work, is in particular often perceived as an admission of weakness, as highlighted in [307]. But we should not despair, if there is one thing to remember from the theory of disordered systems, it is that old things may take a very long time to change, but eventually they always end up doing so. In the meantime, the best we can hope for is to be a modest part of the small fluctuation that will eventually lead the system to a new, but just as temporary, form of consensus.

⁴⁰Our contribution to the boundedly rational description of the Slutsky matrix was for instance judged “more narrow than typical” (when it did not trigger a *bona fide* lecture on the difference between cardinal and ordinal utility functions) by economic journals.

Bibliography

- [1] J. Garnier-Brun, M. Benzaquen, S. Ciliberti, and J.-P. Bouchaud, “A new spin on optimal portfolios and ecological equilibria,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2021, no. 9, p. 093408, 2021.
- [2] J. Garnier-Brun, J.-P. Bouchaud, and M. Benzaquen, “Bounded rationality and animal spirits: a fluctuation-response approach to slusky matrices,” *Journal of Physics: Complexity*, vol. 4, no. 1, p. 015004, 2023.
- [3] J. Garnier-Brun, M. Benzaquen, and J.-P. Bouchaud *In preparation*, 2023.
- [4] A. Dechant, J. Garnier-Brun, and S.-i. Sasa, “Thermodynamic bounds on correlation times,” *arXiv preprint arXiv:2303.13038*, 2023.
- [5] R. Zakine, J. Garnier-Brun, A.-C. Becharat, and M. Benzaquen, “Socioeconomic agents as active matter in nonequilibrium sakoda-schelling models,” *arXiv preprint arXiv:2307.14270*, 2023.
- [6] I. Alemany, J. N. Rose, J. Garnier-Brun, A. D. Scott, and D. J. Doorly, “Random walk diffusion simulations in semi-permeable layered media with varying diffusivity,” *Scientific Reports*, vol. 12, no. 1, p. 10759, 2022.
- [7] D. Bernoulli, “Exposition of a new theory on the measurement of risk (english translation, 1954),” *Econometrica*, vol. 22, pp. 23–36, 1738.
- [8] L. Walras, “Economie et mécanique,” 1909.
- [9] M. H. Stanley, L. A. Amaral, S. V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M. A. Salinger, and H. E. Stanley, “Scaling behaviour in the growth of companies,” *Nature*, vol. 379, no. 6568, pp. 804–806, 1996.
- [10] G. Caldarelli, M. Marsili, and Y.-C. Zhang, “A prototype model of stock exchange,” *Europhysics Letters*, vol. 40, no. 5, p. 479, 1997.

Bibliography

- [11] J.-P. Bouchaud and M. Mézard, “Wealth condensation in a simple model of economy,” *Physica A: Statistical Mechanics and its Applications*, vol. 282, no. 3-4, pp. 536–545, 2000.
- [12] R. Cont and J.-P. Bouchaud, “Herd behavior and aggregate fluctuations in financial markets,” *Macroeconomic dynamics*, vol. 4, no. 2, pp. 170–196, 2000.
- [13] “Theory of rumour spreading in complex social networks,” *Physica A: Statistical Mechanics and its Applications*, vol. 374, no. 1, pp. 457–470, 2007.
- [14] S. Grauwin, E. Bertin, R. Lemoy, and P. Jensen, “Competition between collective and individual dynamics,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 49, pp. 20622–20626, 2009.
- [15] J.-P. Bouchaud, “Crises and Collective Socio-Economic Phenomena: Simple Models and Challenges,” *Journal of Statistical Physics*, vol. 151, pp. 567–606, May 2013.
- [16] X. Gabaix, “Power laws in economics and finance,” *Annu. Rev. Econ.*, vol. 1, no. 1, pp. 255–294, 2009.
- [17] J.-P. Bouchaud and M. Potters, *Theory of financial risk and derivative pricing: from statistical physics to risk management*. Cambridge university press, 2003.
- [18] J.-P. Bouchaud, J. Bonart, J. Donier, and M. Gould, *Trades, quotes and prices: financial markets under the microscope*. Cambridge University Press, 2018.
- [19] D. Colander, H. Föllmer, A. Haas, M. D. Goldberg, K. Juselius, A. Kirman, T. Lux, and B. Sloth, “The financial crisis and the systemic failure of academic economics,” *Univ. of Copenhagen Dept. of Economics Discussion Paper*, no. 09-03, 2009.
- [20] J. D. Farmer and D. Foley, “The economy needs agent-based modelling,” *Nature*, vol. 460, no. 7256, pp. 685–686, 2009.
- [21] J. E. Stiglitz, “Where modern macroeconomics went wrong,” *Oxford Review of Economic Policy*, vol. 34, no. 1-2, pp. 70–106, 2018.
- [22] O. Blanchard, “On the future of macroeconomic models,” *Oxford Review of Economic Policy*, vol. 34, no. 1-2, pp. 43–54, 2018.

- [23] A. G. Haldane and A. E. Turrell, “An interdisciplinary model for macroeconomics,” *Oxford Review of Economic Policy*, vol. 34, no. 1-2, pp. 219–251, 2018.
- [24] J.-P. Bouchaud, “Econophysics: still fringe after 30 years?,” *Europhysics News*, vol. 50, no. 1, pp. 24–27, 2019.
- [25] V. Cannella and J. A. Mydosh, “Magnetic ordering in gold-iron alloys,” *Physical Review B*, vol. 6, no. 11, p. 4220, 1972.
- [26] C. Mulder, A. Van Duyneveldt, and J. Mydosh, “Susceptibility of the Cu Mn spin-glass: Frequency and field dependences,” *Physical Review B*, vol. 23, no. 3, p. 1384, 1981.
- [27] S. F. Edwards and P. W. Anderson, “Theory of spin glasses,” *Journal of Physics F: Metal Physics*, vol. 5, no. 5, p. 965, 1975.
- [28] D. Sherrington and S. Kirkpatrick, “Solvable model of a spin-glass,” *Physical review letters*, vol. 35, no. 26, p. 1792, 1975.
- [29] P. W. Anderson, “Spin glass VI: Spin glass as cornucopia,” *Physics Today*, vol. 42, no. 9, pp. 9–11, 1989.
- [30] Y. Fu and P. W. Anderson, “Application of statistical mechanics to NP-complete problems in combinatorial optimisation,” *Journal of Physics A: Mathematical and General*, vol. 19, no. 9, p. 1605, 1986.
- [31] S. Kirkpatrick and G. Toulouse, “Configuration space analysis of travelling salesman problems,” *Journal de Physique*, vol. 46, no. 8, pp. 1277–1292, 1985.
- [32] M. Mézard and G. Parisi, “A replica analysis of the travelling salesman problem,” *Journal de physique*, vol. 47, no. 8, pp. 1285–1296, 1986.
- [33] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, “Optimization by simulated annealing,” *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [34] W. A. Little, “The existence of persistent states in the brain,” *Mathematical biosciences*, vol. 19, no. 1-2, pp. 101–120, 1974.
- [35] J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities,” *Proceedings of the national academy of sciences*, vol. 79, no. 8, pp. 2554–2558, 1982.

Bibliography

- [36] E. Gardner, B. Derrida, and P. Mottishaw, “Zero temperature parallel dynamics for infinite range spin glasses and neural networks,” *Journal de physique*, vol. 48, no. 5, pp. 741–755, 1987.
- [37] H. Ramsauer, B. Schäfl, J. Lehner, P. Seidl, M. Widrich, T. Adler, L. Gruber, M. Holzleitner, M. Pavlović, G. K. Sandve, *et al.*, “Hopfield networks is all you need,” *arXiv preprint arXiv:2008.02217*, 2020.
- [38] G. Biroli and M. Mézard, “Generative diffusion in very large dimensions,” *arXiv preprint arXiv:2306.03518*, 2023.
- [39] S. A. Kauffman, *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA, 1993.
- [40] H. Rieger, “Solvable model of a complex ecosystem with randomly interacting species,” *Journal of Physics A: Mathematical and General*, vol. 22, no. 17, p. 3447, 1989.
- [41] G. Bunin, “Ecological communities with Lotka-Volterra dynamics,” *Physical Review E*, vol. 95, no. 4, p. 042414, 2017.
- [42] S. Galluccio, J. P. Bouchaud, and M. Potters, “Rational decisions, random matrices and spin glasses,” *Physica A: Statistical Mechanics and its Applications*, vol. 259, no. 3-4, pp. 449–456, 1998.
- [43] D. Challet, M. Marsili, and R. Zecchina, “Statistical mechanics of systems with heterogeneous agents: Minority games,” *Physical Review Letters*, vol. 84, no. 8, p. 1824, 2000.
- [44] M. Mézard, G. Parisi, and M. Virasoro, *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, vol. 9. World Scientific Publishing Company, 1987.
- [45] P. Charbonneau, E. Marinari, G. Parisi, F. Ricci-terseghi, G. Sicuro, F. Zamponi, and M. Mezard, *Spin Glass Theory and Far Beyond: Replica Symmetry Breaking after 40 Years*. World Scientific, 2023.
- [46] R. Lucas, “In defence of the dismal science,” *The Economist*, vol. 6, no. 8, p. 2009, 2009.
- [47] M. Gallegati and A. Kirman, *Beyond the representative agent*. Edward Elgar Publishing, 1999.
- [48] A. Kirman, “The economic crisis is a crisis for economic theory,” *CESifo Economic Studies*, vol. 56, no. 4, pp. 498–535, 2010.

- [49] A. Mas-Colell, M. D. Whinston, J. R. Green, *et al.*, *Microeconomic theory*, vol. 1. Oxford university press New York, 1995.
- [50] I. Fisher, *Mathematical Investigations in the Theory of Value and Prices*. PhD thesis, Yale University, 1891.
- [51] G. Debreu *et al.*, “Representation of a preference ordering by a numerical function,” *Decision processes*, vol. 3, pp. 159–165, 1954.
- [52] G. Gigerenzer and P. M. Todd, *Simple heuristics that make us smart*. Oxford University Press, USA, 1999.
- [53] X. Gabaix, “A sparsity-based model of bounded rationality,” *The Quarterly Journal of Economics*, vol. 129, no. 4, pp. 1661–1710, 2014.
- [54] R. D. Luce, *Individual choice behavior*. Wesley, 1959.
- [55] S. P. Anderson, A. de Palma, and J.-F. Thisse, *Discrete Choice Theory of Product Differentiation*. The MIT Press, 1992.
- [56] D. McFadden, “Conditional logit analysis of qualitative choice behavior,” tech. rep., Institute of Urban and Regional Development, University of California, 1973.
- [57] D. L. McFadden, “Quantal choice analysis: A survey,” *Annals of Economic and Social Measurement, Volume 5, number 4*, pp. 363–390, 1976.
- [58] D. Challet, M. Marsili, and Y.-C. Zhang, *Minority games: interacting agents in financial markets*. OUP Oxford, 2004.
- [59] M. Marsili, “On the multinomial logit model,” *Physica A: Statistical Mechanics and its Applications*, vol. 269, no. 1, pp. 9–15, 1999.
- [60] J.-P. Nadal, G. Weisbuch, O. Chenevez, and A. Kirman, “A formal approach to market organization: choice functions, mean field approximation and maximum entropy principle,” *Advances in Self-Organization and Evolutionary Economics*, pp. 149–159, 1998.
- [61] R. D. Luce, “The choice axiom after twenty years,” *Journal of mathematical psychology*, vol. 15, no. 3, pp. 215–233, 1977.
- [62] J. Rieskamp, J. R. Busemeyer, and B. A. Mellers, “Extending the bounds of rationality: Evidence and theories of preferential choice,” *Journal of Economic Literature*, vol. 44, no. 3, pp. 631–661, 2006.

Bibliography

- [63] H. A. Simon, “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99–118, 1955.
- [64] H. A. Simon, “Rational choice and the structure of the environment.,” *Psychological review*, vol. 63, no. 2, p. 129, 1956.
- [65] B. Mério, A. Borsos, Z. Hosszú, Z. Oláh, and N. Vágó, “A high resolution agent-based model of the hungarian housing market,” tech. rep., MNB Working Papers, 2022.
- [66] S. Gualdi, M. Tarzia, F. Zamponi, and J.-P. Bouchaud, “Tipping points in macroeconomic agent-based models,” *Journal of Economic Dynamics and Control*, vol. 50, pp. 29–61, 2015.
- [67] J. M. Sakoda, “The checkerboard model of social interaction,” *The Journal of Mathematical Sociology*, vol. 1, pp. 119–132, Jan. 1971.
- [68] R. Hegselmann, “Thomas c. schelling and james m. sakoda: The intellectual, technical, and social history of a model,” *Journal of Artificial Societies and Social Simulation*, vol. 20, no. 3, 2017.
- [69] T. C. Schelling, “Dynamic models of segregation,” *Journal of mathematical sociology*, vol. 1, no. 2, pp. 143–186, 1971.
- [70] L. P. Boustan, “Racial residential segregation in american cities,” tech. rep., National Bureau of Economic Research, 2013.
- [71] J. Trounstine, *Segregation by design: Local politics and inequality in American cities*. Cambridge University Press, 2018.
- [72] P. Jensen, T. Matreux, J. Cambe, H. Larralde, and E. Bertin, “Giant Catalytic Effect of Altruists in Schelling’s Segregation Model,” *Physical Review Letters*, vol. 120, p. 208301, May 2018.
- [73] T. Schneider and E. Pytte, “Random-field instability of the ferromagnetic state,” *Physical Review B*, vol. 15, no. 3, p. 1519, 1977.
- [74] Y. Imry and S.-K. Ma, “Random-field instability of the ordered state of continuous symmetry,” *Physical Review Letters*, vol. 35, no. 21, p. 1399, 1975.
- [75] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, “Hysteresis and hierarchies: Dynamics of disorder-driven first-order phase transformations,” *Physical Review Letters*, vol. 70, no. 21, p. 3347, 1993.

- [76] Q. Michard and J.-P. Bouchaud, “Theory of collective opinion shifts: from smooth trends to abrupt swings,” *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 47, pp. 151–159, 2005.
- [77] N. Goldenfeld, *Lectures on phase transitions and the renormalization group*. CRC Press, 2018.
- [78] J. Sengers, “Universality of critical phenomena in classical fluids,” in *Phase Transitions: Cargèse 1980* (J. Levy, J. Zinn-Justin, M. Levy, and J.-C. Le Guillou, eds.), Springer Science & Business Media, 1982.
- [79] J. Le Guillou and J. Zinn-Justin, “Critical exponents for the n-vector model in three dimensions from field theory,” *Physical Review Letters*, vol. 39, no. 2, p. 95, 1977.
- [80] T. Tao, “E pluribus unum: from complexity, universality,” *Daedalus*, vol. 141, no. 3, pp. 23–34, 2012.
- [81] L. Koralov and Y. G. Sinai, *Theory of probability and random processes*. Springer Science & Business Media, 2007.
- [82] P. W. Anderson, “More is different: Broken symmetry and the nature of the hierarchical structure of science.,” *Science*, vol. 177, no. 4047, pp. 393–396, 1972.
- [83] M. Kardar, *Statistical physics of particles*. Cambridge University Press, 2007.
- [84] D. Tong, *Statistical physics*. University of Cambridge, 2011.
- [85] N. G. Van Kampen, *Stochastic processes in physics and chemistry*, vol. 1. Elsevier, 1992.
- [86] U. Seifert, “Entropy production along a stochastic trajectory and an integral fluctuation theorem,” *Physical review letters*, vol. 95, no. 4, p. 040602, 2005.
- [87] D. Landau and K. Binder, *A guide to Monte Carlo simulations in statistical physics*. Cambridge university press, 2021.
- [88] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, “Equation of state calculations by fast computing machines,” *The journal of chemical physics*, vol. 21, no. 6, pp. 1087–1092, 1953.
- [89] R. J. Glauber, “Time-dependent statistics of the ising model,” *Journal of mathematical physics*, vol. 4, no. 2, pp. 294–307, 1963.

Bibliography

- [90] P. Langevin, “Sur la théorie du mouvement Brownien,” *Compt. Rendus*, vol. 146, pp. 530–533, 1908.
- [91] G. E. Uhlenbeck and L. S. Ornstein, “On the theory of the Brownian motion,” *Physical review*, vol. 36, no. 5, p. 823, 1930.
- [92] H. Risken, *Fokker-planck equation*. Springer, 1996.
- [93] G. A. Pavliotis, *Stochastic processes and applications*. Springer, 2016.
- [94] R. G. Palmer, “Broken ergodicity,” *Advances in Physics*, vol. 31, no. 6, pp. 669–735, 1982.
- [95] F. Ritort, “Glassiness in a model without energy barriers,” *Physical review letters*, vol. 75, no. 6, p. 1190, 1995.
- [96] S. Arrhenius, “Über die reaktionsgeschwindigkeit bei der inversion von rohrzucker durch säuren,” *Zeitschrift für physikalische Chemie*, vol. 4, no. 1, pp. 226–248, 1889.
- [97] J.-P. Bouchaud, “Weak ergodicity breaking and aging in disordered systems,” *Journal de Physique I*, vol. 2, no. 9, pp. 1705–1713, 1992.
- [98] A. Dechant, E. Lutz, D. Kessler, and E. Barkai, “Fluctuations of time averages for langevin dynamics in a binding force field,” *Physical review letters*, vol. 107, no. 24, p. 240603, 2011.
- [99] L. F. Cugliandolo and J. Kurchan, “Weak ergodicity breaking in mean-field spin-glass models,” *Philosophical Magazine B*, vol. 71, no. 4, pp. 501–514, 1995.
- [100] E. J. Hinch, *Perturbation Methods*. Cambridge Texts in Applied Mathematics, Cambridge University Press, 1991.
- [101] S. Kirkpatrick and D. Sherrington, “Infinite-ranged models of spin-glasses,” *Physical Review B*, vol. 17, no. 11, p. 4384, 1978.
- [102] J. R. de Almeida and D. J. Thouless, “Stability of the sherrington-kirkpatrick solution of a spin glass model,” *Journal of Physics A: Mathematical and General*, vol. 11, no. 5, p. 983, 1978.
- [103] D. J. Thouless, P. W. Anderson, and R. G. Palmer, “Solution of ‘Solvable model of a spin glass’,” *Philosophical Magazine*, vol. 35, no. 3, pp. 593–601, 1977.

- [104] A. Bray, H. Sompolinsky, and C. Yu, “On the ‘naive’ mean-field equations for spin glasses,” *Journal of Physics C: Solid State Physics*, vol. 19, no. 32, p. 6389, 1986.
- [105] T. Castellani and A. Cavagna, “Spin-glass theory for pedestrians,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2005, no. 5, pp. 215–266, 2005.
- [106] A. J. Bray and M. A. Moore, “Metastable states in spin glasses,” *Journal of Physics C: Solid State Physics*, vol. 13, no. 19, 1980.
- [107] A. Cavagna, I. Giardinà, G. Parisi, and M. Mézard, “On the formal equivalence of the TAP and thermodynamic methods in the SK model,” *Journal of Physics A: Mathematical and General*, vol. 36, no. 5, p. 1175, 2003.
- [108] J. Kurchan, “Replica trick to calculate means of absolute values: applications to stochastic equations,” *Journal of Physics A: Mathematical and General*, vol. 24, no. 21, p. 4969, 1991.
- [109] T. Aspelmeier, A. Bray, and M. Moore, “Complexity of Ising spin glasses,” *Physical review letters*, vol. 92, no. 8, p. 087203, 2004.
- [110] A. Crisanti, L. Leuzzi, G. Parisi, and T. Rizzo, “Spin-glass complexity,” *Physical review letters*, vol. 92, no. 12, p. 127203, 2004.
- [111] J. L. W. V. Jensen, “Sur les fonctions convexes et les inégalités entre les valeurs moyennes,” *Acta mathematica*, vol. 30, no. 1, pp. 175–193, 1906.
- [112] F. Tanaka and S. Edwards, “Analytic theory of the ground state properties of a spin glass. I. Ising spin glass,” *Journal of Physics F: Metal Physics*, vol. 10, no. 12, p. 2769, 1980.
- [113] R. Monasson, “Structural glass transition and the entropy of the metastable states,” *Physical review letters*, vol. 75, no. 15, p. 2847, 1995.
- [114] G. Parisi, “Glasses, replicas and all that,” in *Nonequilibrium Dynamics in Condensed Matter: Les Houches Session LXXVII, 1-26 July, 2002* (J.-L. Barrat, M. Feigelman, J. Kurchan, and J. Dalibard, eds.), Springer, 2004.
- [115] G. Parisi, “Infinite number of order parameters for spin-glasses,” *Physical Review Letters*, vol. 43, no. 23, p. 1754, 1979.
- [116] G. Parisi, “A sequence of approximated solutions to the SK model for spin glasses,” *Journal of Physics A: Mathematical and General*, vol. 13, no. 4, p. L115, 1980.

Bibliography

- [117] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, “Replica symmetry breaking and the nature of the spin glass phase,” *Journal de Physique*, vol. 45, no. 5, pp. 843–854, 1984.
- [118] H. Rieger, “Nonequilibrium dynamics and aging in the three-dimensional Ising spin-glass model,” *Journal of Physics A: Mathematical and General*, vol. 26, no. 15, p. L615, 1993.
- [119] L. Cugliandolo, J. Kurchan, and F. Ritort, “Evidence of aging in spin-glass mean-field models,” *Physical Review B*, vol. 49, no. 9, p. 6331, 1994.
- [120] L. Berthier and J.-P. Bouchaud, “Geometrical aspects of aging and rejuvenation in the Ising spin glass: A numerical study,” *Physical Review B*, vol. 66, no. 5, p. 054404, 2002.
- [121] J.-P. Bouchaud, L. F. Cugliandolo, J. Kurchan, and M. Mézard, “Out of equilibrium dynamics in spin-glasses and other glassy systems,” *Spin glasses and random fields*, vol. 12, p. 161, 1998.
- [122] M. Bernaschi, A. Billoire, A. Maiorano, G. Parisi, and F. Ricci-Tersenghi, “Strong ergodicity breaking in aging of mean-field spin glasses,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 30, pp. 17522–17527, 2020.
- [123] L. F. Cugliandolo and D. S. Dean, “Full dynamical solution for a spherical spin-glass model,” *Journal of Physics A: Mathematical and General*, vol. 28, no. 15, p. 4213, 1995.
- [124] J. M. Kosterlitz, D. J. Thouless, and R. C. Jones, “Spherical model of a spin-glass,” *Physical Review Letters*, vol. 36, no. 20, p. 1217, 1976.
- [125] L. Cugliandolo, J. Kurchan, G. Parisi, and F. Ritort, “Matrix models as solvable glass models,” *Physical review letters*, vol. 74, no. 6, p. 1012, 1995.
- [126] H. Markowitz, “Portfolio selection,” *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
- [127] S. Ciliberti and M. Mézard, “Risk minimization through portfolio replication,” *The European Physical Journal B*, vol. 57, no. 2, pp. 175–180, 2007.
- [128] I. Varga-Haszonits, F. Caccioli, and I. Kondor, “Replica approach to mean-variance portfolio optimization,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2016, no. 12, pp. 1–21, 2016.

- [129] I. Kondor, G. Papp, and F. Caccioli, “Analytic solution to variance optimization with no short positions,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2017, no. 12, pp. 1–29, 2017.
- [130] S. Pafka and I. Kondor, “Noisy covariance matrices and portfolio optimization,” *European Physical Journal B*, vol. 27, no. 2, pp. 277–280, 2002.
- [131] S. Pafka and I. Kondor, “Estimated correlation matrices and portfolio optimization,” *Physica A: Statistical Mechanics and its Applications*, vol. 343, no. 1-4, pp. 623–634, 2004.
- [132] J. Bun, J.-P. Bouchaud, and M. Potters, “Cleaning large correlation matrices: tools from random matrix theory,” *Physics Reports*, vol. 666, pp. 1–109, 2017.
- [133] R. Clarke, H. De Silva, and S. Thorley, “Minimum-variance portfolio composition,” *The Journal of Portfolio Management*, vol. 37, no. 2, pp. 31–45, 2011.
- [134] C.-A. Lehalle and G. Simon, “Portfolio selection with active strategies: how long only constraints shape convictions,” *Journal of Asset Management*, pp. 1–21, 2021.
- [135] P.-A. Reigner, V. Nguyen, S. Ciliberti, P. Seager, and J.-P. Bouchaud, “Agnostic allocation portfolios: A sweet spot in the risk-based jungle?,” *The Journal of Portfolio Management*, vol. 46, no. 4, pp. 22–38, 2020.
- [136] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters, “Noise dressing of financial correlation matrices,” *Physical review letters*, vol. 83, no. 7, p. 1467, 1999.
- [137] W. W. Hager, “Updating the inverse of a matrix,” *SIAM review*, vol. 31, no. 2, pp. 221–239, 1989.
- [138] R. M. May, “Will a large complex system be stable?,” *Nature*, vol. 238, no. 5364, pp. 413–414, 1972.
- [139] F. Courchamp, T. Clutton-Brock, and B. Grenfell, “Inverse density dependence and the Allee effect,” *Trends in ecology & evolution*, vol. 14, no. 10, pp. 405–410, 1999.
- [140] J. K. Sankaran and A. A. Patil, “On the optimal selection of portfolios under limited diversification,” *Journal of banking & Finance*, vol. 23, no. 11, pp. 1655–1666, 1999.

Bibliography

- [141] R. Courant and D. Hilbert, *Methods of Mathematical Physics*, vol. 2: Partial Differential Equations. Wiley-VCH, 1962.
- [142] V. Azcoiti, E. Follana, and F. Ritort, “Static chaos in spin glasses: the case of quenched disorder perturbations,” *Journal of Physics A: Mathematical and General*, vol. 28, no. 14, p. 3863, 1995.
- [143] F. Krzakala and J. P. Bouchaud, “Disorder chaos in spin glasses,” *Europhysics Letters*, vol. 72, no. 3, pp. 472–478, 2005.
- [144] C. Monthus and T. Garel, “Chaos properties of the one-dimensional long-range ising spin-glass,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2014, no. 3, p. P03020, 2014.
- [145] I. Kondor, “On chaos in spin glasses,” *Journal of Physics A: Mathematical and General*, vol. 22, no. 5, p. L163, 1989.
- [146] G. Biroli, G. Bunin, and C. Cammarota, “Marginally stable equilibria in critical ecosystems,” *New Journal of Physics*, vol. 20, no. 8, p. 083051, 2018.
- [147] V. Ros, F. Roy, G. Biroli, G. Bunin, and A. M. Turner, “Generalized lotka-volterra equations with random, nonreciprocal interactions: The typical number of equilibria,” *Physical Review Letters*, vol. 130, no. 25, p. 257401, 2023.
- [148] V. Ros, F. Roy, G. Biroli, and G. Bunin, “Quenched complexity of equilibria for asymmetric generalized lotka-volterra equations,” *Journal of Physics A: Mathematical and Theoretical*.
- [149] J. Moran and J.-P. Bouchaud, “May’s instability in large economies,” *Phys. Rev. E*, vol. 100, p. 032307, Sep 2019.
- [150] T. Dessertaine, J. Moran, M. Benzaquen, and J.-P. Bouchaud, “Out-of-equilibrium dynamics and excess volatility in firm networks,” *Journal of Economic Dynamics and Control*, vol. 138, p. 104362, 2022.
- [151] A. Mas-Colell, M. D. Whinston, J. R. Green, *et al.*, *Microeconomic theory*, vol. 1. Oxford university press New York, 1995.
- [152] J. Elster and I. Ekeland, *Théorie économique et rationalité*. Vuibert, 2011.
- [153] U. M. B. Marconi, A. Puglisi, L. Rondoni, and A. Vulpiani, “Fluctuation–dissipation: response theory in statistical physics,” *Physics reports*, vol. 461, no. 4-6, pp. 111–195, 2008.
- [154] T. Veblen, *The theory of the leisure class*. 1899.

- [155] F. M. Bass, “A new product growth for model consumer durables,” *Management science*, vol. 15, no. 5, pp. 215–227, 1969.
- [156] W. A. Brock and S. N. Durlauf, “Discrete choice with social interactions,” *The Review of Economic Studies*, vol. 68, no. 2, pp. 235–260, 2001.
- [157] M. J. Salganik, P. S. Dodds, and D. J. Watts, “Experimental study of inequality and unpredictability in an artificial cultural market,” *Science*, vol. 311, no. 5762, pp. 854–856, 2006.
- [158] M. B. Gordon, J.-P. Nadal, D. Phan, and V. Semeshenko, “Discrete choices under social influence: Generic properties,” *Mathematical Models and Methods in Applied Sciences*, vol. 19, no. supp01, pp. 1441–1481, 2009.
- [159] J. Y. Campbell and J. H. Cochrane, “By force of habit: A consumption-based explanation of aggregate stock market behavior,” *Journal of Political Economy*, vol. 107, no. 2, pp. 205–251, 1999.
- [160] A. al Nowaihi and L. Stracca, “Keeping up with the Joneses, reference dependence, and equilibrium indeterminacy,” Working Paper Series 444, European Central Bank, Feb. 2005.
- [161] F. G. Morelli, M. Benzaquen, M. Tarzia, and J.-P. Bouchaud, “Confidence collapse in a multihousehold, self-reflexive dsge model,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 17, pp. 9244–9249, 2020.
- [162] J. M. Keynes, “The general theory of employment,” *The quarterly journal of economics*, vol. 51, no. 2, pp. 209–223, 1937.
- [163] C. Borghesi and J.-P. Bouchaud, “Of songs and men: a model for multiple choice with herding,” *Quality & Quantity*, vol. 41, pp. 557–568, Mar. 2007.
- [164] A. Lucas, “Non-equilibrium phase transitions in competitive markets caused by network effects,” *arXiv preprint arXiv:2204.05314*, 2022.
- [165] M. Browning and P.-A. Chiappori, “Efficient intra-household allocations: A general characterization and empirical tests,” *Econometrica*, p. 1241–1278, 1998.
- [166] P. Lecomte, *Texture-induced hydrophobicity, and some collective effects in granular matter and economics*. PhD thesis, Institut Polytechnique de Paris, 2019.
- [167] J.-M. Grandmont, “Expectations formation and stability of large socioeconomic systems,” *Econometrica*, pp. 741–781, 1998.

Bibliography

- [168] A. Kirman, *Complex economics: individual and collective rationality*. Routledge, 2010.
- [169] G. Dosi and A. Roventini, “More is different... and complex! the case for agent-based macroeconomics,” *Journal of Evolutionary Economics*, vol. 29, pp. 1–37, 2019.
- [170] M. King and J. Kay, *Radical uncertainty: Decision-making for an unknowable future*. Hachette UK, 2020.
- [171] A. Marcet and T. J. Sargent, “Convergence of least squares learning mechanisms in self-referential linear stochastic models,” *Journal of Economic theory*, vol. 48, no. 2, pp. 337–368, 1989.
- [172] G. W. Evans, S. Honkapohja, *et al.*, “Learning as a rational foundation for macroeconomics and finance,” *Rethinking expectations: The way forward for macroeconomics*, vol. 68, 2013.
- [173] R. Pemantle, “A survey of random processes with reinforcement,” *Probability Surveys*, vol. 4, pp. 1–79, 2007.
- [174] D. Lamberton, P. Tarrès, *et al.*, “When can the two-armed bandit algorithm be trusted?,” *Ann. Appl. Probab.*, vol. 14, no. 1, pp. 1424–1454, 2004.
- [175] J. Moran, A. Fosset, D. Luzzati, J.-P. Bouchaud, and M. Benzaquen, “By force of habit: Self-trapping in a dynamical utility landscape,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 30, no. 5, p. 053123, 2020.
- [176] C. Colon and J.-P. Bouchaud, “The radical complexity of rewiring supplier–buyer networks,” *Available at SSRN 4300311*, 2022.
- [177] T. Galla and J. D. Farmer, “Complex dynamics in learning complicated games,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 4, pp. 1232–1236, 2013.
- [178] J.-P. Bouchaud and R. E. Farmer, “Self-fulfilling prophecies, quasi nonergodicity, and wealth inequality,” *Journal of Political Economy*, vol. 131, no. 4, pp. 000–000, 2023.
- [179] G. Parisi, “Complex systems: a physicist’s viewpoint,” *Physica A: Statistical Mechanics and its Applications*, vol. 263, no. 1, pp. 557–564, 1999. Proceedings of the 20th IUPAP International Conference on Statistical Physics.
- [180] G. Parisi, “Physics complexity and biology,” *Advances in Complex Systems*, vol. 10, no. supp02, pp. 223–232, 2007.

- [181] J.-P. Bouchaud, “Radical complexity,” *Entropy*, vol. 23, no. 12, p. 1676, 2021.
- [182] D. Sharma, J.-P. Bouchaud, M. Tarzia, and F. Zamponi, “Good speciation and endogenous business cycles in a constraint satisfaction macroeconomic model,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2021, no. 6, p. 063403, 2021.
- [183] C. J. Watkins and P. Dayan, “Q-learning,” *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [184] C. Camerer and T. Hua Ho, “Experience-weighted attraction learning in normal form games,” *Econometrica*, vol. 67, no. 4, pp. 827–874, 1999.
- [185] D. Panchenko, “Introduction to the sk model,” *Current Developments in Mathematics*, pp. 231–291, 2015.
- [186] Y. Sato, E. Akiyama, and J. P. Crutchfield, “Stability and diversity in collective adaptation,” *Physica D: Nonlinear Phenomena*, vol. 210, no. 1-2, pp. 21–57, 2005.
- [187] T. Galla, “Intrinsic noise in game dynamical learning,” *Physical review letters*, vol. 103, no. 19, p. 198702, 2009.
- [188] T. Galla, “Cycles of cooperation and defection in imperfect learning,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2011, no. 08, p. P08007, 2011.
- [189] D. Vilone, A. Robledo, and A. Sánchez, “Chaos and unpredictability in evolutionary dynamics in discrete time,” *Physical review letters*, vol. 107, no. 3, p. 038101, 2011.
- [190] A. Kianercy and A. Galstyan, “Dynamics of Boltzmann Q learning in two-player two-action games,” *Physical Review E*, vol. 85, no. 4, p. 041145, 2012.
- [191] J. Burridge, “Limit cycles and the benefits of a short memory in rock-paper-scissors games,” *Physical Review E*, vol. 92, no. 4, p. 042111, 2015.
- [192] J. B. Sanders, J. D. Farmer, and T. Galla, “The prevalence of chaotic dynamics in games with many players,” *Scientific reports*, vol. 8, no. 1, pp. 1–13, 2018.
- [193] V. Semeshenko, M. B. Gordon, J.-P. Nadal, and D. Phan, “Choice under social influence: effects of learning behaviours on the collective dynamics,” *Contributions to Economic Analysis*, vol. 280, pp. 177–203, 2006.

Bibliography

- [194] V. Semeshenko, M. B. Gordon, and J.-P. Nadal, “Collective states in social systems with interacting learning agents,” *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 19-20, pp. 4903–4916, 2008.
- [195] M. Oppen and S. Diederich, “Phase transition and $1/f$ noise in a game dynamical model,” *Physical review letters*, vol. 69, no. 10, p. 1616, 1992.
- [196] T. Galla, “Random replicators with asymmetric couplings,” *Journal of Physics A: Mathematical and General*, vol. 39, no. 15, p. 3853, 2006.
- [197] A. Leonidov, A. Savvateev, and A. G. Semenov, “Ising game on graphs,” *arXiv preprint arXiv:2108.00824*, 2021.
- [198] A. Leonidov and E. Vasilyeva, “Strategic stiffening/cooling in the ising game,” *Chaos, Solitons & Fractals*, vol. 160, p. 112279, 2022.
- [199] A. Montanari, “Optimization of the sherrington–kirkpatrick hamiltonian,” *SIAM Journal on Computing*, no. 0, pp. FOCS19–1, 2021.
- [200] S. Boettcher, “Simulations of ground state fluctuations in mean-field ising spin glasses,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2010, no. 07, p. P07002, 2010.
- [201] S. Hwang, V. Folli, E. Lanza, G. Parisi, G. Ruocco, and F. Zamponi, “On the number of limit cycles in asymmetric neural networks,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2019, no. 5, p. 053402, 2019.
- [202] H.-J. Sommers and W. Dupont, “Distribution of frozen fields in the mean-field theory of spin glasses,” *Journal of Physics C: Solid State Physics*, vol. 17, no. 32, p. 5785, 1984.
- [203] S. Pankov, “Low-temperature solution of the sherrington-kirkpatrick model,” *Physical review letters*, vol. 96, no. 19, p. 197204, 2006.
- [204] M. Müller and M. Wyart, “Marginal stability in structural, spin, and electron glasses,” *Annu. Rev. Condens. Matter Phys.*, vol. 6, no. 1, pp. 177–200, 2015.
- [205] P. Cizeau and J.-P. Bouchaud, “Mean field theory of dilute spin-glasses with power-law interactions,” *Journal of Physics A: Mathematical and General*, vol. 26, no. 5, p. L187, 1993.
- [206] K. Janzen, A. Engel, and M. Mézard, “Thermodynamics of the lévy spin glass,” *Physical Review E*, vol. 82, no. 2, p. 021127, 2010.

- [207] I. Neri, F. Metz, and D. Bollé, “The phase diagram of lévy spin glasses,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2010, no. 01, p. P01010, 2010.
- [208] T. Aspelmeier, “Bond chaos in the Sherrington–Kirkpatrick model,” *Journal of Physics A: Mathematical and Theoretical*, vol. 41, no. 20, p. 205005, 2008.
- [209] G. Toulouse, “On the mean field theory of mixed spin glass-ferromagnetic phases,” *Journal de Physique Lettres*, vol. 41, no. 18, pp. 447–449, 1980.
- [210] R. Zecchina, “Liquid computing in neural networks.” Towards a theory of artificial and biological neural networks – Les Houches workshop, 2023.
- [211] H. Takayama and K. Nemoto, “Spin glass properties of a class of mean-field models,” *Journal of Physics: Condensed Matter*, vol. 2, no. 8, p. 1997, 1990.
- [212] K. Nishimura, K. Nemoto, and H. Takayama, “Metastable states of the naive mean-field model for spin glasses at finite temperatures,” *Journal of Physics A: Mathematical and General*, vol. 23, no. 24, p. 5915, 1990.
- [213] H. J. Sommers, A. Crisanti, H. Sompolinsky, and Y. Stein, “Spectrum of large random asymmetric matrices,” *Physical review letters*, vol. 60, no. 19, p. 1895, 1988.
- [214] G. Parisi, “Computing the number of metastable states in infinite-range models,” in *Mathematical Statistical Physics: Lecture Notes of the Les Houches Summer School 2005* (A. Bovier, F. Dunlop, A. Van Enter, F. Den Hollander, and J. Dalibard, eds.), Elsevier, 2006.
- [215] F. Waugh, C. Marcus, and R. Westervelt, “Fixed-point attractors in analog neural computation,” *Physical review letters*, vol. 64, no. 16, p. 1986, 1990.
- [216] H. Gutfreund, J. Reger, and A. Young, “The nature of attractors in an asymmetric spin glass with deterministic dynamics,” *Journal of Physics A: Mathematical and General*, vol. 21, no. 12, p. 2775, 1988.
- [217] U. Bastolla and G. Parisi, “Relaxation, closing probabilities and transition from oscillatory to chaotic attractors in asymmetric neural networks,” *Journal of Physics A: Mathematical and General*, vol. 31, no. 20, p. 4583, 1998.
- [218] S. Franz and G. Parisi, “Recipes for metastable states in spin glasses,” *Journal de Physique I*, vol. 5, no. 11, pp. 1401–1415, 1995.
- [219] L. F. Cugliandolo, “Recent applications of dynamical mean-field methods,” *arXiv preprint arXiv:2305.01229*, 2023.

Bibliography

- [220] H. Eissfeller and M. Opper, “New method for studying the dynamics of disordered spin systems without finite-size effects,” *Physical review letters*, vol. 68, no. 13, p. 2094, 1992.
- [221] F. Roy, G. Biroli, G. Bunin, and C. Cammarota, “Numerical implementation of dynamical mean field theory for disordered systems: Application to the Lotka–Volterra model of ecosystems,” *Journal of Physics A: Mathematical and Theoretical*, vol. 52, no. 48, p. 484001, 2019.
- [222] H. Eissfeller and M. Opper, “Mean-field Monte Carlo approach to the Sherrington-Kirkpatrick model with asymmetric couplings,” *Physical Review E*, vol. 50, no. 2, p. 709, 1994.
- [223] T. Dessertaine, *Large macroeconomic fluctuations: self-organized criticality in firm networks, Agent Based Models and random matrices*. PhD thesis, Institut Polytechnique de Paris, 2022.
- [224] H. Sompolinsky, A. Crisanti, and H.-J. Sommers, “Chaos in random neural networks,” *Physical review letters*, vol. 61, no. 3, p. 259, 1988.
- [225] J. Heimerl and A. De Martino, “Broken ergodicity and memory in the minority game,” *Journal of Physics A: Mathematical and General*, vol. 34, no. 40, p. L539, 2001.
- [226] A. Altieri, F. Roy, C. Cammarota, and G. Biroli, “Properties of equilibria and glassy phases of the random Lotka-Volterra model with demographic noise,” *Physical Review Letters*, vol. 126, no. 25, p. 258301, 2021.
- [227] D. Martí, N. Brunel, and S. Ostojic, “Correlations between synapses in pairs of neurons slow down dynamics in randomly connected neural networks,” *Physical Review E*, vol. 97, no. 6, p. 062314, 2018.
- [228] H. Yoshino, “Off-equilibrium dynamics of a $(1+1)$ -dimensional directed polymer in random media,” *Journal of Physics A: Mathematical and General*, vol. 29, no. 7, p. 1421, 1996.
- [229] T. A. de Pirey and G. Bunin, “Aging by near-extinctions in many-variable interacting populations,” *Physical Review Letters*, vol. 130, no. 9, p. 098401, 2023.
- [230] L. F. Cugliandolo, J. Kurchan, P. Le Doussal, and L. Peliti, “Glassy behaviour in disordered systems with nonrelaxational dynamics,” *Physical review letters*, vol. 78, no. 2, p. 350, 1997.

- [231] L. Berthier and J. Kurchan, “Non-equilibrium glass transitions in driven and active matter,” *Nature Physics*, vol. 9, no. 5, pp. 310–314, 2013.
- [232] E. Marinari, G. Parisi, F. Ricci-Tersenghi, and J. J. Ruiz-Lorenzo, “Off-equilibrium dynamics at very low temperatures in three-dimensional spin glasses,” *Journal of Physics A: Mathematical and General*, vol. 33, no. 12, p. 2373, 2000.
- [233] L. Molgedey, J. Schuchhardt, and H. G. Schuster, “Suppressing chaos in neural networks by noise,” *Physical review letters*, vol. 69, no. 26, p. 3717, 1992.
- [234] F. Aguirre-Lopez *In preparation*, 2023.
- [235] A. Antonov, A. Leonidov, and A. Semenov, “Self-excited ising game,” *Physica A: Statistical Mechanics and its Applications*, vol. 561, p. 125305, 2021.
- [236] C. Van Vreeswijk and H. Sompolinsky, “Chaos in neuronal networks with balanced excitatory and inhibitory activity,” *Science*, vol. 274, no. 5293, pp. 1724–1726, 1996.
- [237] N. Brunel, “Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons,” *Journal of computational neuroscience*, vol. 8, pp. 183–208, 2000.
- [238] K. Rajan, L. Abbott, and H. Sompolinsky, “Stimulus-dependent suppression of chaos in recurrent neural networks,” *Physical review e*, vol. 82, no. 1, p. 011903, 2010.
- [239] M. Stern, H. Sompolinsky, and L. F. Abbott, “Dynamics of random neural networks with bistable units,” *Physical Review E*, vol. 90, no. 6, p. 062710, 2014.
- [240] J. Kadmon and H. Sompolinsky, “Transition to chaos in random neuronal networks,” *Physical Review X*, vol. 5, no. 4, p. 041030, 2015.
- [241] D. G. Clark and L. Abbott, “Theory of coupled neuronal-synaptic dynamics,” *arXiv preprint arXiv:2302.08985*, 2023.
- [242] U. Pereira-Obilinovic, J. Aljadeff, and N. Brunel, “Forgetting leads to chaos in attractor networks,” *Physical Review X*, vol. 13, no. 1, p. 011009, 2023.
- [243] A. P. Kirman and N. J. Vriend, “Learning to be loyal. a study of the marseille fish market,” in *Interaction and market structure: Essays on heterogeneity in economics*, pp. 33–56, Springer, 2000.

Bibliography

- [244] T. Aspelmeier, R. Blythe, A. J. Bray, and M. A. Moore, “Free-energy landscapes, dynamics, and the edge of chaos in mean-field models of spin glasses,” *Physical Review B*, vol. 74, no. 18, p. 184411, 2006.
- [245] T. Aspelmeier and M. Moore, “Realizable solutions of the Thouless-Anderson-Palmer equations,” *Physical Review E*, vol. 100, no. 3, p. 032127, 2019.
- [246] D. Vinković and A. Kirman, “A physical analogue of the Schelling model,” *Proceedings of the National Academy of Sciences*, vol. 103, pp. 19261–19265, Dec. 2006. Publisher: Proceedings of the National Academy of Sciences.
- [247] L. Dall’Asta, C. Castellano, and M. Marsili, “Statistical physics of the Schelling model of segregation,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, p. L07002, July 2008.
- [248] L. Gauvin, J. Vannimenus, and J.-P. Nadal, “Phase diagram of a Schelling segregation model,” *The European Physical Journal B*, vol. 70, pp. 293–304, July 2009.
- [249] T. Rogers and A. J. McKane, “A unified framework for schelling’s model of segregation,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2011, no. 07, p. P07006, 2011.
- [250] L. Gauvin, J.-P. Nadal, and J. Vannimenus, “Schelling segregation in an open city: A kinetically constrained Blume-Emery-Griffiths spin-1 system,” *Physical Review E*, vol. 81, p. 066120, June 2010.
- [251] G. Barmpalias, R. Elwes, and A. Lewis-Pye, “Tipping points in 1-dimensional schelling models with switching agents,” *Journal of Statistical Physics*, vol. 158, pp. 806–852, 2015.
- [252] D. Ortega, J. Rodríguez-Laguna, and E. Korutcheva, “A Schelling model with a variable threshold in a closed city segregation model. analysis of the universality classes,” *Physica A: Statistical Mechanics and its Applications*, vol. 574, p. 126010, 2021.
- [253] D. Ortega, J. Rodríguez-Laguna, and E. Korutcheva, “Avalanches in an extended schelling model: An explanation of urban gentrification,” *Physica A: Statistical Mechanics and its Applications*, vol. 573, p. 125943, 2021.
- [254] D. Abella, M. San Miguel, and J. J. Ramasco, “Aging effects in schelling segregation model,” *Scientific Reports*, vol. 12, no. 1, p. 19376, 2022.

- [255] A. Kolmogoroff, “Zur theorie der markoffschen ketten,” *Mathematische Annalen*, vol. 112, no. 1, pp. 155–160, 1936.
- [256] P. Tamayo and W. Klein, “Critical dynamics and global conservation laws,” *Physical Review Letters*, vol. 63, pp. 2757–2759, Dec. 1989.
- [257] A. J. Bray, “Comment on “Critical dynamics and global conservation laws’ ”,” *Physical Review Letters*, vol. 66, pp. 2048–2048, Apr. 1991.
- [258] A. D. Rutenberg, “Nonequilibrium phase ordering with a global conservation law,” *Physical Review E*, vol. 54, pp. 972–973, July 1996.
- [259] A. Tartaglia, L. F. Cugliandolo, and M. Picco, “Coarsening and percolation in the kinetic 2 d Ising model with spin exchange updates and the voter model,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2018, p. 083202, Aug. 2018.
- [260] J. T. Siebert, F. Dittrich, F. Schmid, K. Binder, T. Speck, and P. Virnau, “Critical behavior of active Brownian particles,” *Physical Review E*, vol. 98, p. 030601, Sept. 2018. arXiv:1712.02258 [cond-mat].
- [261] C. Maggi, M. Paoluzzi, A. Crisanti, E. Zaccarelli, and N. Gnan, “Universality class of the motility-induced critical point in large scale off-lattice simulations of active particles,” *Soft Matter*, vol. 17, no. 14, pp. 3807–3812, 2021.
- [262] F. Dittrich, T. Speck, and P. Virnau, “Critical behavior in active lattice models of motility-induced phase separation,” *The European Physical Journal E*, vol. 44, p. 53, Apr. 2021.
- [263] N. Gnan and C. Maggi, “Critical behavior of quorum-sensing active particles,” *Soft Matter*, vol. 18, no. 39, pp. 7654–7661, 2022.
- [264] A. I. Curatolo, N. Zhou, Y. Zhao, C. Liu, A. Daerr, J. Tailleur, and J. Huang, “Cooperative pattern formation in multi-component bacterial systems through reciprocal motility regulation,” *Nature Physics*, vol. 16, pp. 1152–1157, Nov. 2020.
- [265] A. I. Curatolo, M. R. Evans, Y. Kafri, and J. Tailleur, “Multilane driven diffusive systems,” *Journal of Physics A: Mathematical and Theoretical*, vol. 49, p. 095601, Jan. 2016. Publisher: IOP Publishing.
- [266] S. M. Allen and J. W. Cahn, “A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening,” *Acta metallurgica*, vol. 27, no. 6, pp. 1085–1095, 1979.

Bibliography

- [267] J. W. Cahn and J. E. Hilliard, “Free Energy of a Nonuniform System. I. Interfacial Free Energy,” *The Journal of Chemical Physics*, vol. 28, pp. 258–267, 08 2004.
- [268] T. Grafke, M. E. Cates, and E. Vanden-Eijnden, “Spatiotemporal Self-Organization of Fluctuating Bacterial Colonies,” *Physical Review Letters*, vol. 119, p. 188003, Nov. 2017.
- [269] J. O’Byrne and J. Tailleur, “Lamellar to Micellar Phases and Beyond: When Tactic Active Systems Admit Free Energy Functionals,” *Physical Review Letters*, vol. 125, p. 208003, Nov. 2020.
- [270] J. O’Byrne, “Nonequilibrium currents in stochastic field theories: A geometric insight,” *Physical Review E*, vol. 107, p. 054105, May 2023.
- [271] A. Lefevre and G. Biroli, “Dynamics of interacting particle systems: stochastic process and field theory,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2007, no. 07, p. P07024, 2007.
- [272] J. Tailleur, J. Kurchan, and V. Lecomte, “Mapping out-of-equilibrium into equilibrium in one-dimensional transport models,” *Journal of Physics A: Mathematical and Theoretical*, vol. 41, p. 505001, Dec. 2008.
- [273] L. Bertini, A. De Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim, “Macroscopic fluctuation theory,” *Reviews of Modern Physics*, vol. 87, pp. 593–636, June 2015.
- [274] Y. Baek, Y. Kafri, and V. Lecomte, “Dynamical phase transitions in the current distribution of driven diffusive channels,” *Journal of Physics A: Mathematical and Theoretical*, vol. 51, p. 105001, Mar. 2018.
- [275] R. Zakine and E. Vanden-Eijnden, “Minimum action method for nonequilibrium phase transitions,” *arXiv preprint arXiv:2202.06936*, 2022.
- [276] A. P. Solon, J. Stenhammar, M. E. Cates, Y. Kafri, and J. Tailleur, “Generalized thermodynamics of phase equilibria in scalar active matter,” *Phys. Rev. E*, vol. 97, p. 020602, Feb 2018.
- [277] A. P. Solon, J. Stenhammar, M. E. Cates, Y. Kafri, and J. Tailleur, “Generalized thermodynamics of motility-induced phase separation: phase equilibria, Laplace pressure, and change of ensembles,” *New Journal of Physics*, vol. 20, p. 075001, July 2018.
- [278] S. Saha, J. Agudo-Canalejo, and R. Golestanian, “Scalar Active Mixtures: The Nonreciprocal Cahn-Hilliard Model,” *Physical Review X*, vol. 10, p. 041009, Oct. 2020.

- [279] A. Dinelli, J. O’Byrne, A. Curatolo, Y. Zhao, P. Sollich, and J. Tailleur, “Non-reciprocity across scales in active mixtures,” *arXiv preprint arXiv:2203.07757*, 2022.
- [280] A.-C. Becharat *et al.* *In preparation*, 2023.
- [281] C.-R. Hwang, S.-Y. Hwang-Ma, and S.-J. Sheu, “Accelerating diffusions,” 2005.
- [282] A. Ichiki and M. Ohzeki, “Violation of detailed balance accelerates relaxation,” *Physical Review E*, vol. 88, no. 2, p. 020101, 2013.
- [283] C.-R. Hwang, R. Normand, and S.-J. Wu, “Variance reduction for diffusions,” *Stochastic Processes and their Applications*, vol. 125, no. 9, pp. 3522–3540, 2015.
- [284] A. B. Duncan, T. Lelievre, and G. A. Pavliotis, “Variance reduction using nonreversible langevin samplers,” *Journal of statistical physics*, vol. 163, pp. 457–491, 2016.
- [285] F. Bouchet and J. Reygner, “Generalisation of the eyring–kramers transition rate formula to irreversible diffusion processes,” in *Annales Henri Poincaré*, vol. 17, pp. 3499–3532, Springer, 2016.
- [286] D. Borgis and M. Moreau, “On the escape rate from a metastable state in a nonpotential system,” *Physica A: Statistical Mechanics and its Applications*, vol. 163, no. 3, pp. 877–894, 1990.
- [287] S. Bradde and G. Biroli, “The generalized arrhenius law in out of equilibrium systems,” *arXiv preprint arXiv:1204.6027*, 2012.
- [288] A. C. Barato and U. Seifert, “Thermodynamic uncertainty relation for biomolecular processes,” *Physical review letters*, vol. 114, no. 15, p. 158101, 2015.
- [289] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, “Dissipation bounds all steady-state current fluctuations,” *Physical review letters*, vol. 116, no. 12, p. 120601, 2016.
- [290] T. Nemoto and S.-i. Sasa, “Thermodynamic formula for the cumulant generating function of time-averaged current,” *Physical Review E*, vol. 84, no. 6, p. 061113, 2011.
- [291] A. Dechant and S.-i. Sasa, “Current fluctuations and transport efficiency for general langevin systems,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2018, no. 6, p. 063209, 2018.

Bibliography

- [292] J. O’Byrne, “Nonequilibrium currents in stochastic field theories: A geometric insight,” *Physical Review E*, vol. 107, no. 5, p. 054105, 2023.
- [293] A. Dechant, “Thermodynamic constraints on the power spectral density in and out of equilibrium,” *arXiv preprint arXiv:2306.00417*, 2023.
- [294] M. Ohzeki and A. Ichiki, “Langevin dynamics neglecting detailed balance condition,” *Physical Review E*, vol. 92, no. 1, p. 012105, 2015.
- [295] F. Ghimenti and F. van Wijland, “Accelerating, to some extent, the p -spin dynamics,” *Physical Review E*, vol. 105, no. 5, p. 054137, 2022.
- [296] A. Dechant *et al.* *In preparation*, 2023.
- [297] R. L. Stratonovich, “On the probability functional of diffusion processes,” *Selected Trans. in Math. Stat. Prob.*, vol. 10, pp. 273–286, 1971.
- [298] D. Dürr and A. Bach, “The onsager-machlup function as lagrangian for the most probable path of a diffusion process,” *Communications in Mathematical Physics*, vol. 60, pp. 153–170, 1978.
- [299] J. Gladrow, U. F. Keyser, R. Adhikari, and J. Kappler, “Experimental measurement of relative path probabilities and stochastic actions,” *Physical Review X*, vol. 11, no. 3, p. 031022, 2021.
- [300] E. Weinan, W. Ren, and E. Vanden-Eijnden, “Minimum action method for the study of rare events,” *Communications on pure and applied mathematics*, vol. 57, no. 5, pp. 637–656, 2004.
- [301] M. Heymann and E. Vanden-Eijnden, “The geometric minimum action method: A least action principle on the space of curves,” *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, vol. 61, no. 8, pp. 1052–1117, 2008.
- [302] T. Grafke and E. Vanden-Eijnden, “Numerical computation of rare events via large deviation theory,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, no. 6, 2019.
- [303] L. Kikuchi, R. Adhikari, and J. Kappler, “Diffusivity dependence of the transition path ensemble,” *arXiv preprint arXiv:2203.12947*, 2022.
- [304] T. Grafke and A. Laio, “Metadynamics for transition paths in irreversible dynamics,” *arXiv preprint arXiv:2211.09476*, 2022.

- [305] D. S. Dean, I. Drummond, and R. Horgan, “Effective transport properties for diffusion in random media,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2007, no. 07, p. P07013, 2007.
- [306] C. Touya and D. Dean, “Dynamical transition for a particle in a squared gaussian potential,” *Journal of Physics A: Mathematical and Theoretical*, vol. 40, no. 5, p. 919, 2007.
- [307] J.-P. Bouchaud, “De la physique statistique aux sciences sociales : les défis de la pluridisciplinarité,” Collège de France, Leçon inaugurale de la chaire annuelle Innovation technologique, 2020.
- [308] J. Spanier and K. B. Oldham, *An atlas of functions*, ch. 42, pp. 405–410. Hemisphere publishing corporation New York, 1987.
- [309] V. H. Moll, *Special Integrals of Gradshteyn and Ryzhik: the Proofs*, vol. 1. CRC Press, 2015.
- [310] M. E. Newman and G. T. Barkema, *Monte Carlo methods in statistical physics*. Clarendon Press, 1999.
- [311] P. Glasserman and Y.-C. Ho, *Gradient estimation via perturbation analysis*, vol. 116. Springer Science & Business Media, 1991.
- [312] P. Glasserman, *Monte Carlo methods in financial engineering*, vol. 53. Springer Science & Business Media, 2013.
- [313] A. Javanmard, A. Montanari, and F. Ricci-Tersenghi, “Phase transitions in semidefinite relaxations,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 16, pp. E2218–E2223, 2016.

Bibliography

Appendices

Appendix A

Optimal portfolios

This section reproduces the Appendices of [1].

A.1 Full self-consistent equation

Starting from

$$\tilde{\beta}^+ = \frac{\sum_j \beta_j^2 \theta_j / z_j + 1}{\sum_j \beta_j \theta_j / z_j} \quad (\text{A.1})$$

we make use of the central limit theorem as for $N \gg 1$

$$\frac{1}{N} \sum_j \beta_j^k \theta_j / z_j \simeq \langle z^{-1} \rangle \langle \beta^k \rangle + \frac{1}{\sqrt{N}} \xi_k \quad (\text{A.2})$$

where ξ_k are Gaussian noises with mean $\langle \xi_k \rangle = 0$ and variance

$$\langle \xi_k^2 \rangle = \langle z^{-2} \rangle \langle \beta^{2k} \rangle - \langle z^{-1} \rangle^2 \langle \beta^k \rangle^2. \quad (\text{A.3})$$

After factorisation and expansion of the denominator, Eq. (A.1) can be written as

$$\begin{aligned} \tilde{\beta}^+ = & \frac{\langle \beta^2 \rangle}{\langle \beta \rangle} + \frac{1}{N} \left[\frac{\bar{z}}{\langle \beta \rangle} - \frac{\bar{z}^2}{\langle \beta \rangle^2} \left(\langle \xi_1 \xi_2 \rangle - \frac{\langle \beta^2 \rangle}{\langle \beta \rangle} \langle \xi_1^2 \rangle \right) \right] + O\left(\frac{1}{N^2}\right) \\ & + \underbrace{\frac{1}{\sqrt{N}} \frac{\bar{z}}{\langle \beta \rangle} \left(\xi_2 - \frac{\langle \beta^2 \rangle}{\langle \beta \rangle} \xi_1 \right)}_{\text{fluctuations}} + O\left(\frac{1}{N^{3/2}}\right), \end{aligned} \quad (\text{A.4})$$

with $\bar{z} = \langle z^{-1} \rangle^{-1}$, which can be rewritten as $\tilde{\beta}^+ = \beta^+ + \frac{1}{\sqrt{N}} \xi$ with the final noise term

$$\xi = \frac{\bar{z}}{\langle \beta \rangle} \left(\xi_2 - \frac{\langle \beta^2 \rangle}{\langle \beta \rangle} \xi_1 \right) \quad (\text{A.5})$$

Appendix A. Optimal portfolios

that still has zero mean and variance

$$\gamma^2 = \frac{\bar{z}^2 \langle z^{-2} \rangle}{\langle \beta \rangle^2} \left(\langle \beta^4 \rangle - 2 \frac{\langle \beta^2 \rangle \langle \beta^3 \rangle}{\langle \beta \rangle} + \left(\frac{\langle \beta^2 \rangle}{\langle \beta \rangle} \right)^2 \langle \beta^2 \rangle \right). \quad (\text{A.6})$$

Substituting the correct values for $\langle \xi_1 \xi_2 \rangle$ and $\langle \xi_1^2 \rangle$, the deterministic term can be rewritten as

$$\beta^+ = \frac{\langle \beta^2 \rangle}{\langle \beta \rangle} + \frac{1}{N} \left[\frac{\bar{z}}{\langle \beta \rangle} - \frac{\bar{z}^2 \langle z^{-2} \rangle}{\langle \beta \rangle^2} \left(\langle \beta^3 \rangle - \frac{\langle \beta^2 \rangle^2}{\langle \beta \rangle} \right) \right] + O\left(\frac{1}{N^2}\right). \quad (\text{A.7})$$

Now, as detailed in the following section for the case $\alpha = 1$, at the leading order one may Taylor expand the averages about the threshold β^+ as $F_N(\beta) = H(N^\alpha(\beta - \beta^+))$, with $\alpha \geq 1/2$ from the form of F_N found in the main text. As a result, $\langle \beta^k \rangle \simeq m + \sigma(\dots) + \dots + \sigma^k(\dots)$ and so in the regime $\sigma = \chi/N$ we have both

$$\langle \beta^3 \rangle - \frac{\langle \beta^2 \rangle^2}{\langle \beta \rangle} = 0 + O\left(\frac{1}{N}\right) \quad \text{and} \quad \langle \beta^4 \rangle - 2 \frac{\langle \beta^2 \rangle \langle \beta^3 \rangle}{\langle \beta \rangle} + \left(\frac{\langle \beta^2 \rangle}{\langle \beta \rangle} \right)^2 \langle \beta^2 \rangle = 0 + O\left(\frac{1}{N}\right).$$

As such, the second term in N^{-1} may be eliminated from Eq. (A.7), and we find that the variance γ^2 is at most of order N^{-1} .

A.2 Sommerfeld-like expansions

Starting from the averages

$$\langle \beta^k \rangle = \int_{-\infty}^{\infty} d\beta \beta^k \rho(\beta) F_N(\beta), \quad (\text{A.8})$$

we may define $\psi(\beta) = \int_{-\infty}^{\beta} d\beta \beta^k \rho(\beta)$ such that integrating by parts

$$\int_{-\infty}^{\infty} d\beta \beta^k \rho(\beta) F_N(\beta) = - \int_{-\infty}^{\infty} d\beta \psi(\beta) F'_N(\beta). \quad (\text{A.9})$$

as the boundary term vanishes given $\psi(-\infty)\rho(-\infty) = 0$ and $F_N(+\infty) = 0$. As previously mentioned, $F'_N(\beta)$ is peaked in a small region around β^+ , therefore we can expand ψ in this region with a Taylor series, giving in turn

$$\begin{aligned} \int_{-\infty}^{\infty} d\beta \varphi(\beta) F_N(\beta) &= - \psi(\beta^+) \int_{-\infty}^{\infty} d\beta F'_N(\beta) \\ &\quad - \psi'(\beta^+) \int_{-\infty}^{\infty} d\beta (\beta - \beta^+) F'_N(\beta) + O((\beta - \beta^+)^2). \end{aligned} \quad (\text{A.10})$$

Now, changing variables as $x = N(\beta - \beta^+)$ the equation becomes

$$\int_{-\infty}^{\infty} d\beta \varphi(\beta) F_N(\beta) = -\psi(\beta^+) \int_{-\infty}^{\infty} dx H'(x) - \frac{1}{N} \psi'(\beta^+) \int_{-\infty}^{\infty} dx x H'(x) + O(N^{-2}) \quad (\text{A.11})$$

where the first integral in x easily gives -1 given our knowledge of $H(x)$, while the second can be written as

$$\kappa = - \int_{-\infty}^{\infty} dx x H'(x) = \int_0^{\infty} dx (H(x) + H(-x) - 1) \quad (\text{A.12})$$

that will clearly be zero in the case of a symmetric function written like $H(x > 0) = \frac{1}{2} + \epsilon(x)$ and $H(x < 0) = \frac{1}{2} - \epsilon(x)$. Given that in our case $H(x)$ is a complementary error function and is thus symmetric, we therefore have

$$\int_{-\infty}^{\infty} d\beta \beta^k \rho(\beta) F_N(\beta) = \int_{-\infty}^{\beta^+} d\beta \beta^k \rho(\beta) + O\left(\frac{1}{N^2}\right) \quad (\text{A.13})$$

or, using the compact notations introduced in the main text,

$$\langle\langle \beta^k \rangle\rangle = \langle \beta^k \rangle_c + O\left(\frac{1}{N^2}\right) \quad (\text{A.14})$$

A.3 Reaction term

We essentially adapt the Onsager cavity field approach to our problem. Considering a system with N assets and their associated spins $\{\theta\}$, the threshold for inclusion was shown to be given by

$$\tilde{\beta}_N^+ = \frac{\sum_j \beta_j^2 \theta_j / z_j + 1}{\sum_j \beta_j \theta_j / z_j} = \beta_N^+ + \frac{1}{\sqrt{N}} \xi \quad (\text{A.15})$$

After introduction of a new asset, at the index 0 for simplicity, this threshold is altered as

$$\tilde{\beta}_{N+1}^+ = \frac{\sum_j \beta_j^2 \theta_j / z_j + \beta_0^2 \theta_0 / z_0 + 1}{\sum_j \beta_j \theta_j / z_j + \beta_0 \theta_0 / z_0} \quad (\text{A.16})$$

that can be expressed, after applying the central limit theorem to sums and expanding the denominator as before, as

$$\tilde{\beta}_{N+1}^+ = \beta_N^+ + \frac{1}{\sqrt{N}} \xi + \frac{1}{N} c(\beta_0) \quad (\text{A.17})$$

with the reaction term

$$c(\beta_0) = \frac{1}{\langle\langle \beta \rangle\rangle} \left(\frac{\beta_0^2 \theta_0}{z_0} - \frac{\langle\langle \beta^2 \rangle\rangle}{\langle\langle \beta \rangle\rangle} \frac{\beta_0 \theta_0}{z_0} \right). \quad (\text{A.18})$$

Appendix A. Optimal portfolios

Now, just like we took $F_N(\beta) = \text{Prob}(\xi \geq \beta - \beta^+)$, we have

$$F_{N+1}(\beta) = \text{Prob}\left(\xi \geq \beta - \beta^+ - \frac{1}{N}c(\beta_0)\right),$$

which may be Taylor expanded and averaged over the distribution β_0 such that

$$F_{N+1}(\beta) = F_N(\beta) + \frac{1}{N} \frac{e^{-\frac{1}{2}\left(\frac{\beta - \beta^+}{\gamma}\right)^2}}{\sqrt{2\pi\gamma^2}} \int_{-\infty}^{\infty} d\beta_0 c(\beta_0) \rho(\beta_0) F_N(\beta_0) + O\left(\frac{1}{N^2}\right). \quad (\text{A.19})$$

Then, simply going back to the definition of the averages $\langle\langle \beta^k \rangle\rangle$, we clearly find the integral in the second term becomes

$$\langle\langle c(\beta_0) \rangle\rangle = \frac{1}{\langle\langle \beta \rangle\rangle z_0} \left(\langle\langle \beta^2 \rangle\rangle - \frac{\langle\langle \beta^2 \rangle\rangle}{\langle\langle \beta \rangle\rangle} \langle\langle \beta \rangle\rangle \right) = 0. \quad (\text{A.20})$$

As such, the reaction term has no contribution at order $1/N$, and the naive self-consistent equation for β^+ requires no further modification.

A.4 Detailed resolution of the characteristic equations

Starting with the Gaussian case, the characteristic equation for x , rewritten as

$$s = \int_0^x dv \frac{e^{-v\varphi'(\sigma v)}}{\varphi(\sigma v)} \quad (\text{A.21})$$

is split between the small σv region, that yields a constant contribution, and the large σv region where we had the asymptote

$$\varphi(\sigma v) = \frac{\sqrt{2 \log \sigma v}}{\sigma v}. \quad (\text{A.22})$$

Using this expression, we explicitly write the derivative in the exponent

$$\varphi'(\sigma v) = \frac{1 - 2 \log \sigma v}{\sigma v^2 \sqrt{2 \log \sigma v}}. \quad (\text{A.23})$$

Now, for $v \gg 1$, $v\varphi'(\sigma v)$ decreases like $\sqrt{\log \sigma v}/(\sigma v) \gg 1$, justifying a Taylor expansion of the exponential. As such,

$$e^{-v\varphi'(\sigma v)} = 1 + \frac{2 \log \sigma v - 1}{\sigma v \sqrt{2 \log \sigma v}} + O\left(\frac{\log \sigma v}{(\sigma v)^2}\right) \quad (\text{A.24})$$

and thus the integrand of Eq. (A.21) is given by

$$\frac{e^{-v\varphi'(\sigma v)}}{\varphi(\sigma v)} = \frac{\sigma v}{\sqrt{2\log \sigma v}} + 1 - \frac{1}{2\log \sigma v} + O\left(\frac{\sqrt{\log \sigma v}}{\sigma v}\right). \quad (\text{A.25})$$

The integration of the first term presents a slight challenge, but taking the change of variable $\sigma v = e^{\frac{w^2}{2}}$,

$$\int^x dv \frac{\sigma v}{\sqrt{2\log \sigma v}} = \frac{1}{\sigma} \int^{\sqrt{2\log \sigma x}} dw e^{w^2} = \frac{1}{\sigma} e^{2\log \sigma x} F\left(\sqrt{2\log \sigma x}\right) \quad (\text{A.26})$$

with F the Dawson integral function. Using the asymptotic expansion of this special function [308], this first term finally becomes

$$\frac{\sigma x^2}{2\sqrt{2\log \sigma x}} \left[1 + O\left(\frac{1}{\log \sigma x}\right)\right]. \quad (\text{A.27})$$

The third term in Eq. (A.25) can also be expanded asymptotically, as

$$\frac{1}{2} \int^x dv \frac{1}{\log \sigma v} = \frac{\text{li}(\sigma x)}{2\sigma} = \frac{x}{2\log \sigma x} \left[1 + O\left(\frac{1}{\log \sigma x}\right)\right] \quad (\text{A.28})$$

where equalities are up to an additive constant, and li is the well known logarithmic integral function. Bringing everything together,

$$s = \frac{\sigma x^2}{2\sqrt{2\log \sigma x}} \left[1 + O\left(\frac{1}{\log \sigma x}\right)\right] + x \left[1 - \frac{1}{2\log \sigma x} + O\left(\frac{1}{(\log \sigma x)^2}\right)\right] + O\left((\log \sigma x)^{\frac{3}{2}}\right). \quad (\text{A.29})$$

In the large x limit, the very first term will largely dominate others and we may therefore recover the expression given in the main text,

$$s \sim \frac{\sigma x^2}{2\sqrt{2\log \sigma x}}. \quad (\text{A.30})$$

Rearranging this expression, we have

$$\sigma x = \sqrt{2\sigma s}(2\log \sigma x)^{\frac{1}{4}} \quad (\text{A.31})$$

and thus we can take the iterated logarithm

$$\log \sigma x = \frac{1}{2} \log \sigma s + \frac{3}{4} \log 2 + \frac{1}{4} \log \log \sigma x \quad (\text{A.32})$$

and thus

$$\sigma x = \sqrt{2\sigma s}(\log \sigma s)^{\frac{1}{4}} \left[1 + O\left(\frac{\log \log \sigma s}{\log \sigma s}\right)\right]^{\frac{1}{4}}, \quad (\text{A.33})$$

Appendix A. Optimal portfolios

giving the asymptotic relation

$$x(s) \sim \sqrt{\frac{2s}{\sigma}} (\log \sigma s)^{\frac{1}{4}}. \quad (\text{A.34})$$

The final characteristic ODE can then be integrated,

$$\log z = \int_0^s dv e^{x(v)\varphi'(\sigma x(v))} \varphi(\sigma x(v)), \quad (\text{A.35})$$

once again splitting the constant contribution from the small σv region and the known asymptotic behaviour. Replacing with the expression for $x(s)$, we have

$$\log \sigma x(v) = \frac{1}{2} \log \sigma v \left[1 + O\left(\frac{\log \log \sigma v}{\log \sigma v}\right) \right] \quad (\text{A.36})$$

and thus

$$x(v)\varphi'(\sigma x(v)) = -\frac{(\log \sigma v)^{\frac{1}{4}}}{\sqrt{\sigma v}} \left[1 + O\left(\frac{\log \log \sigma v - 1}{\log \sigma v}\right) \right]. \quad (\text{A.37})$$

For $\sigma v \gg 1$, the exponential term in Eq. (A.35) can therefore be Taylor expanded. Given

$$\varphi(\sigma x(v)) = \frac{(\log \sigma v)^{\frac{1}{4}}}{\sqrt{2\sigma v}} \left[1 + O\left(\frac{\log \log \sigma v}{\log \sigma v}\right) \right], \quad (\text{A.38})$$

the integral can finally be written as

$$\log z = \int^s dv \frac{(\log \sigma v)^{\frac{1}{4}}}{\sqrt{2\sigma v}} - \int^s dv \frac{\sqrt{\log \sigma v}}{\sqrt{2\sigma v}} + O\left(\int^s dv \frac{\log \log \sigma v}{\sqrt{\sigma v}(\log \sigma v)^{\frac{3}{4}}}\right). \quad (\text{A.39})$$

Now, the first term will clearly dominate for large s . The integral may be evaluating by taking the change of variable $\sigma v = e^{2w}$:

$$\int^s dv \frac{(\log \sigma v)^{\frac{1}{4}}}{\sqrt{2\sigma v}} = \frac{2^{\frac{3}{4}}}{\sigma} \int^{\frac{1}{2} \log \sigma s} dw w^{\frac{1}{4}} e^w = \sqrt{\frac{2s}{\sigma}} (\log \sigma s)^{\frac{1}{4}} \left[1 + O\left(\frac{1}{\log \sigma s}\right) \right] \quad (\text{A.40})$$

where the final equality may be shown by integrating by parts [309].

So far, we have used the position along the characteristic s , however integrating the first characteristic equation with the associated boundary condition, we have $s = t$. Recalling that the continuous variable t is analogous to the size of the problem N and z to the average number of solutions $\langle \mathcal{N}_s \rangle$, we may finally express the result with the quantites of interest

$$\langle \mathcal{N}_s \rangle \sim \exp \left\{ \sqrt{\frac{2N}{\sigma}} (\log \sigma N)^{\frac{1}{4}} \right\}, \quad (\text{A.41})$$

as given in the main text. It should be noted that as the error terms are of logarithmic orders, we expect the convergence to this asymptote to be relatively slow in N .

For the uniform case, the calculations are much easier thanks to the simpler form of the maximum sparsity. For $\sigma v > 2^{-\frac{1}{2}}$ large we remind that

$$\varphi(\sigma v) = \frac{1}{2^{\frac{1}{4}}\sqrt{\sigma v}}, \quad (\text{A.42})$$

giving in turn

$$v\varphi'(\sigma v) = -\frac{1}{2^{\frac{5}{4}}\sqrt{\sigma v}}. \quad (\text{A.43})$$

The integral given in Eq. (A.21) therefore amounts to

$$s = \int^s dv \left[2^{\frac{1}{4}}\sqrt{\sigma v} + \frac{1}{2} + O\left(\frac{1}{\sqrt{\sigma v}}\right) \right] \quad (\text{A.44})$$

after Taylor expanding the exponential term, and removing the constant contribution by using the boundary condition. One then easily finds

$$s = \frac{2^{\frac{5}{4}}\sqrt{\sigma x^{\frac{3}{2}}}}{3} \left[1 + O\left(\frac{1}{\sqrt{\sigma x}}\right) \right], \quad (\text{A.45})$$

and thus for $\sigma x \gg 1$, which is expected for large N ,

$$x(s) \sim \left(\frac{3}{2} \frac{s}{\sqrt{\sqrt{2}\sigma}} \right)^{\frac{2}{3}}. \quad (\text{A.46})$$

As for the Gaussian case, this may be reinjected in the expressions of φ and φ' to calculate $\log z$. We find

$$\varphi(\sigma x(v)) = \left(\frac{2}{3} \frac{1}{\sqrt{2}\sigma v} \right)^{\frac{1}{3}} \quad \text{and} \quad x(v)\varphi'(\sigma x(v)) = -\frac{1}{2} \left(\frac{2}{3} \frac{1}{\sqrt{2}\sigma v} \right)^{\frac{1}{3}}, \quad (\text{A.47})$$

resulting in, after Taylor expanding the exponential term,

$$\begin{aligned} \log z &= \int^s dv \left(\frac{2}{3} \frac{1}{\sqrt{2}\sigma s} \right)^{\frac{1}{3}} + \frac{1}{2} \int^s dv \left(\frac{2}{3} \frac{1}{\sqrt{2}\sigma s} \right)^{\frac{2}{3}} + O\left(\int^s dv \frac{1}{\sigma v} \right) \\ &= \left(\frac{3}{2} \frac{s}{\sqrt{\sqrt{2}\sigma}} \right)^{\frac{2}{3}} \left[1 + O\left(\frac{1}{\sqrt[3]{\sigma s}} \right) \right]. \end{aligned} \quad (\text{A.48})$$

Like before, realising $s = t$ directly gives the mean number of solutions and associated annealed complexity as a function of N .

A.5 Generalized normal distribution

Taking the generalised normal distribution and performing the change of variable $u = \frac{\beta-1}{\sqrt{2}}$, we have

$$m = \frac{b}{2\sigma\Gamma(1/b)} \int_{-\infty}^{\frac{\beta^+-1}{\sqrt{2}}} du e^{-(\frac{|u|}{\sigma})^b}, \quad (\text{A.49})$$

$$\langle \beta \rangle_c = m + \frac{b\sqrt{2}}{2\sigma\Gamma(1/b)} \int_{-\infty}^{\frac{\beta^+-1}{\sqrt{2}}} du u e^{-(\frac{|u|}{\sigma})^b} \quad (\text{A.50})$$

and finally

$$\begin{aligned} \langle \beta^2 \rangle_c &= m + \frac{b\sqrt{2}}{\sigma\Gamma(1/b)} \int_{-\infty}^{\frac{\beta^+-1}{\sqrt{2}}} du u e^{-(\frac{|u|}{\sigma})^b} \\ &\quad + \frac{b}{\sigma\Gamma(1/b)} \int_{-\infty}^{\frac{\beta^+-1}{\sqrt{2}}} du u^2 e^{-(\frac{|u|}{\sigma})^b}. \end{aligned} \quad (\text{A.51})$$

At this stage, one may first notice that if $\sigma \sim N^{-1}$, then the integrals involving u^k will be of order N^{-k} . As such, as we are interested only in terms in N^{-1} or higher, rewriting $m = \varphi$ and $\langle \beta \rangle_c = \varphi - \psi$ we therefore also have $\langle \beta^2 \rangle_c = \varphi - 2\psi + O(N^{-2})$, where ψ scales as N^{-1} . As such, the self-consistent equation simplifies to

$$\beta^+ = \frac{\varphi - 2\psi}{\varphi - \psi} + \frac{1}{N} \frac{\bar{z}}{\varphi} + O\left(\frac{1}{N^2}\right). \quad (\text{A.52})$$

At this stage, we can reintroduce the ansatz $\beta^+ = 1 + \chi f(\chi)/N$ such that at order N^{-1} the self-consistent equation becomes

$$\chi f(\chi) = \frac{\bar{z}}{\varphi} - \frac{N\psi}{\varphi} \quad (\text{A.53})$$

that corresponds to the equation given in the main text

$$\chi f(\chi) = \frac{\bar{z}}{m} - \frac{\chi}{m} \frac{b}{2\sqrt{2}\Gamma(1/b)} \int_{-f(\chi)}^{\infty} du u e^{-(\frac{|u|}{\sqrt{2}})^b}. \quad (\text{A.54})$$

We now look at the case of finite b . The known asymptote

$$\int_a^b dt f(t) e^{xt} \sim e^{xb} \left[\sum_{k=1}^n (-1)^{k-1} f^{(k-1)}(b) x^{-k} \right] \quad (\text{A.55})$$

as $x \rightarrow \infty$ can then be used with minor tweaking. For all three integrals, we can use the fact that in the regime of interest $\beta^+ - 1 < 0$ so being careful with signs

one may take the substitution $t = u^b$ to have an integrand in the form of the formula above. This then yields

$$\varphi \simeq \frac{1}{2\Gamma(1/b)} e^{-\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^b} \left[\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{1-b} - \frac{b-1}{b} \left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{1-2b} \right] \quad (\text{A.56})$$

$$\psi \simeq \frac{\chi}{N\sqrt{2}\Gamma(1/b)} e^{-\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^b} \left[\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{2-b} - \frac{b-2}{b} \left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{2-2b} \right] \quad (\text{A.57})$$

giving, once plugged in Eq. (A.53)

$$\begin{aligned} 2\Gamma(1/b) e^{\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^b} = & \chi f(\chi) \left[\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{1-b} - \frac{b-1}{b} \left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{1-2b} \right] \\ & + \sqrt{2}\chi \left[\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{2-b} - \frac{b-2}{b} \left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{2-2b} \right]. \end{aligned} \quad (\text{A.58})$$

Finally, given the we know that $\beta^+ < 1$ in this regime, then we can simply realise that $f(\chi) = -|f(\chi)|$ and so terms on the right hand side cancel out to give

$$e^{\left(\frac{|f(\chi)|}{\sqrt{2}}\right)^b} = \frac{\chi}{\sqrt{2}b\Gamma(1/b)} \left(\frac{|f(\chi)|}{\sqrt{2}}\right)^{2-2b}, \quad (\text{A.59})$$

the final equation to be approximated.

A.6 Heterogeneous returns and growth rates

To extend our model to more realistic conditions, it is important to generalise the calculations to variable values of μ_i . In such a case, the long-only portfolio and population dynamics calculations are no longer strictly equivalent, as in the portfolio μ_i and z_i are independent while regarding species we have $z_i = \mu_i/k_i$.

Starting with the long-only portfolio, we now have

$$\theta_i = \Theta \left(\mu_i - \beta_i \frac{\sum_j \beta_j \mu_j / z_j}{\sum_j \beta_j^2 / z_j + 1} \right). \quad (\text{A.60})$$

The first important aspect to notice here is that, just as was the case for β , we may fix the mean of μ to 1 without loss of generality as any multiplicative rescaling $\mu_i \rightarrow \alpha \mu_i$ leaves the above equation invariant. From there, we may proceed as before, now defining a threshold on the quantity $\phi = \beta/\mu$:

$$\tilde{\phi}^+ = \frac{\sum_j \phi_j^2 \mu_j^2 \theta_j / z_j + 1}{\sum_j \phi_j \mu_j^2 \theta_j / z_j}. \quad (\text{A.61})$$

Appendix A. Optimal portfolios

The probability for an asset to be included in the portfolio is now given by the function $F_N(\phi)$ that is qualitatively identical to the previously introduced equivalent for β alone. Taking μ_i and β_i to have the joint probability distribution $\rho(\beta, \mu)$, we must calculate

$$\langle\langle \phi^k \mu^2 \rangle\rangle = \int d\phi \int d\mu \int d\beta \phi^k F_N(\phi) \mu^2 \rho(\beta, \mu) \delta\left(\phi - \frac{\beta}{\mu}\right). \quad (\text{A.62})$$

From the results of A.2, we know that this integral can be calculated as

$$\langle\langle \phi^k \mu^2 \rangle\rangle = \int_{-\infty}^{\phi^+} d\phi \phi^k h(\phi) + O\left(\frac{1}{N^2}\right) \quad (\text{A.63})$$

where the key step is therefore calculating

$$h(\phi) = \int d\beta \int d\mu \mu^2 \rho(\beta, \mu) \delta\left(\phi - \frac{\beta}{\mu}\right) \quad (\text{A.64})$$

that will act as an effective distribution of ϕ . Assuming now that both β and μ are distributed around 1 with a standard deviation scaling in N^{-1} , we may shift and rescale the problem by taking

$$\beta = 1 + \frac{x}{N}, \quad \mu = 1 + \frac{y}{N}, \quad \phi = 1 + \frac{w}{N} \quad (\text{A.65})$$

and $\tilde{\rho}(x, y)$ that is the distribution of β and μ centred now at 0 and with standard deviation of order 1. Using the scaling property of the Dirac delta distribution, we finally have the rescaled effective density for ϕ that is given by

$$\begin{aligned} \tilde{h}(u) &= \int dx \int dy \left(1 + \frac{2y}{N}\right) \tilde{\rho}(x, y) \delta(w - x + y) + O\left(\frac{1}{N^2}\right) \\ &= \int dy \tilde{\rho}(u + y, y) + O\left(\frac{1}{N}\right), \end{aligned} \quad (\text{A.66})$$

as it will be shortly apparent that the N^{-1} contribution vanishes in the self-consistent equation. Note that when $\tilde{\rho}(x, y)$ is a bivariate Gaussian, $\tilde{h}(u)$ is also Gaussian.

Expressing the threshold in ϕ^+ as

$$\phi^+ = 1 + \frac{\tilde{f}}{N}, \quad (\text{A.67})$$

where $\tilde{f}$ depends on the parameters describing the distribution $\tilde{h}$. We may go back to the self-consistent equation that is analogous to that for β^+ and reads

$$\phi^+ = \frac{\langle \phi^2 \rangle_c}{\langle \phi \rangle_c} + \frac{1}{N} \frac{1}{\langle \phi \rangle_c} + O\left(\frac{1}{N^2}\right) \quad (\text{A.68})$$

with now

$$\langle \phi^k \rangle_c = \int_{-\infty}^{\phi^+} d\phi \phi^k h(\phi) = \int_{-\infty}^{\tilde{f}} du \left(1 + \frac{ku}{N} \right) \tilde{h}(u) + O\left(\frac{1}{N^2}\right). \quad (\text{A.69})$$

As such, we finally obtain an equation that is almost identical to the homogeneous $\mu_i = 1$ case, contributions of order 1 cancel out and we have, at order N^{-1} ,

$$\tilde{f} = \frac{1}{m} + \frac{1}{m} \int_{-\infty}^{\tilde{f}} du u \tilde{h}(u) \quad (\text{A.70})$$

with

$$m = \int_{-\infty}^{\tilde{f}} du \tilde{h}(u). \quad (\text{A.71})$$

Clearly, the N^{-1} term in the expression of $\tilde{h}(u)$ is dominated and therefore the μ^2 term that was initially present turns out to be inconsequential.

With this result in mind, we can look at the equilibrium ecosystem problem. As previously mentioned, the relation $z_i = \mu_i/k_i$ means that the threshold is now given by

$$\tilde{\phi}^+ = \frac{\sum_j \phi_j^2 k_j \mu_j \theta_j + 1}{\sum_j \phi_j k_j \mu_j \theta_j}. \quad (\text{A.72})$$

Thus, leaving aside the k_j that play no part in the inclusion or not of the species (as was the case for z_j), the problem is identical to the long-only portfolio problem, albeit with $\langle \phi^k \mu \rangle$ to be calculated instead of $\langle \phi^k \mu^2 \rangle$. Having just determined that when taking μ to be distributed in an N^{-1} region about its mean its contribution in the integral is negligible, we find ourselves with exactly the same self-consistent equation as above.

Appendix A. Optimal portfolios

Appendix B

Slutsky matrices

This section reproduces the Appendices of [2].

B.1 General “thermodynamic” relations

B.1.1 Single agent Slutsky matrix

Assuming the equilibrium distribution is reached, Eq. (4.6) can be rearranged to be rewritten through correlation functions. Indeed, from the quotient rule the first term becomes

$$\frac{\partial}{\partial p_j} \langle x_i \rangle = \frac{1}{Z} \int d\mathbf{x} x_i x_j e^{\beta u(\mathbf{x})} \delta'(\mathbf{x} \cdot \mathbf{p} - w) - \frac{\langle x_i \rangle}{Z} \int d\mathbf{x} x_j e^{\beta u(\mathbf{x})} \delta'(\mathbf{x} \cdot \mathbf{p} - w). \quad (\text{B.1})$$

Now, using the fact that

$$\delta'(\mathbf{x} \cdot \mathbf{p} - w) = -\partial_w \delta(\mathbf{x} \cdot \mathbf{p} - w), \quad (\text{B.2})$$

then for a test function $f(\mathbf{x})$

$$\frac{1}{Z} \int d\mathbf{x} f(\mathbf{x}) e^{\beta u(\mathbf{x})} \delta'(\mathbf{x} \cdot \mathbf{p} - w) = -\frac{\partial}{\partial w} \langle f(\mathbf{x}) \rangle - \langle f(\mathbf{x}) \rangle \partial_w \log Z. \quad (\text{B.3})$$

Therefore bringing everything together we have

$$S_{ij} = \langle x_i \rangle \partial_w \langle x_j \rangle + \langle x_i \rangle \langle x_j \rangle \partial_w \log Z - \partial_w \langle x_i x_j \rangle - \langle x_i x_j \rangle \partial_w \log Z + \langle x_j \rangle \partial_w \langle x_i \rangle, \quad (\text{B.4})$$

or in a more compact form, with $\langle x_i x_j \rangle_c = \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$ and $\Gamma = \partial_w \log Z$,

$$S_{ij} = -\Gamma \langle x_i x_j \rangle_c - \partial_w \langle x_i x_j \rangle_c. \quad (\text{B.5})$$

Appendix B. Slutsky matrices

B.1.2 Many agent Slutsky matrix

The agent-specific Slutsky matrix is defined as

$$S_{ij}^\alpha = \frac{\partial}{\partial p_j} \langle x_i^\alpha \rangle + \langle x_j^\alpha \rangle \frac{\partial}{\partial w^\alpha} \langle x_i^\alpha \rangle. \quad (\text{B.6})$$

Once again assuming equilibrium is reached, we wish to rewrite this expression through correlations to retrieve some information on the symmetry of the matrix. As before, we look at the first term in the above equation,

$$\begin{aligned} \frac{\partial}{\partial p_j} \langle x_i^\alpha \rangle &= \frac{1}{Z_N} \int \mathcal{D}\mathbf{x} x_i^\alpha e^{\beta U(\{\mathbf{x}^\gamma\})} \frac{\partial}{\partial p_j} \prod_\gamma \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma) \\ &\quad - \frac{\langle x_i^\alpha \rangle}{Z_N} \int \mathcal{D}\mathbf{x} e^{\beta U(\{\mathbf{x}^\gamma\})} \frac{\partial}{\partial p_j} \prod_\gamma \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma), \end{aligned} \quad (\text{B.7})$$

with hereafter the shorthand notation $\int \mathcal{D}\mathbf{x} := \int_+ \prod_{\alpha,i} dx_i^\alpha$. Now, by the generalized chain rule we have

$$\frac{\partial}{\partial p_j} \prod_\gamma \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma) = - \sum_\gamma \frac{\partial}{\partial w^\gamma} x_j^\gamma \prod_\eta \delta(\mathbf{x}^\eta \cdot \mathbf{p} - w^\eta), \quad (\text{B.8})$$

which may be plugged back into Eq. (B.7). Noticing that

$$\left\langle \frac{\partial}{\partial w^\gamma} f(\mathbf{x}^\gamma) \right\rangle = \frac{\partial}{\partial w^\gamma} \langle f(\mathbf{x}^\gamma) \rangle + \Gamma_\gamma \langle f(\mathbf{x}^\gamma) \rangle, \quad (\text{B.9})$$

with $\Gamma_\gamma = \frac{\partial}{\partial w^\gamma} \log Z_N$, then

$$\frac{\partial}{\partial p_j} \langle x_i^\alpha \rangle = - \sum_\gamma \left[\Gamma_\gamma \langle x_i^\alpha x_j^\gamma \rangle_c + \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle - \langle x_i^\alpha \rangle \frac{\partial}{\partial w^\gamma} \langle x_j^\gamma \rangle \right]. \quad (\text{B.10})$$

Bringing everything together, we finally have

$$S_{ij}^\alpha = - \sum_\gamma \left[\Gamma_\gamma \langle x_i^\alpha x_j^\gamma \rangle_c + \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle - \langle x_i^\alpha \rangle \frac{\partial}{\partial w^\gamma} \langle x_j^\gamma \rangle \right] + \langle x_j^\alpha \rangle \frac{\partial}{\partial w^\alpha} \langle x_i^\alpha \rangle. \quad (\text{B.11})$$

When there are no interactions, $\gamma \neq \alpha$, $\langle x_i^\alpha x_j^\gamma \rangle_c = 0$, as well as $\frac{\partial}{\partial w^\gamma} \langle x_i^\alpha \rangle = 0$. As a result, we are simply left with

$$S_{ij}^\alpha = -\Gamma_\alpha \langle x_i^\alpha x_j^\alpha \rangle_c - \frac{\partial}{\partial w^\alpha} \langle x_i^\alpha x_j^\alpha \rangle_c, \quad (\text{B.12})$$

which is symmetric and identical to the single-agent system, as expected.

B.1.3 Fluctuation-response relations

We start from the most general expression for the partition function

$$Z_N = \int \mathcal{D}\mathbf{x} e^{\beta U(\{\mathbf{x}^\gamma\})} \prod_{\gamma} \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma), \quad (\text{B.13})$$

and set-out to find a readily measurable expression for $\Gamma_\alpha = \frac{\partial}{\partial w^\alpha} \log Z_N$. Starting from the computation of

$$\frac{\partial Z_N}{\partial w^\alpha} = \int \mathcal{D}\mathbf{x} e^{\beta U(\{\mathbf{x}^\gamma\})} \frac{\partial}{\partial w^\alpha} \prod_{\gamma} \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma); \quad (\text{B.14})$$

we employ the generalized chain rule,

$$\begin{aligned} \frac{\partial}{\partial w^\alpha} \prod_{\gamma} \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma) &= \sum_{\gamma} \frac{\partial}{\partial w^\alpha} \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma) \prod_{\eta \neq \gamma} \delta(\mathbf{x}^\eta \cdot \mathbf{p} - w^\eta) \\ &= -\frac{1}{p_i} \sum_{\gamma} \frac{\partial w^\gamma}{\partial w^\alpha} \frac{\partial}{\partial x_i^\gamma} \delta(\mathbf{x}^\gamma \cdot \mathbf{p} - w^\gamma) \prod_{\eta \neq \gamma} \delta(\mathbf{x}^\eta \cdot \mathbf{p} - w^\eta) \quad \forall i. \end{aligned} \quad (\text{B.15})$$

Now, regardless of the assumption on their distribution, the agents' individual budgets are not related, and as such $\frac{\partial w^\gamma}{\partial w^\alpha} = \delta_{\alpha\gamma}$. Reinjecting in the previous expression and integrating by parts, we obtain

$$\frac{\partial Z_N}{\partial w^\alpha} = \frac{1}{p_i} \int \mathcal{D}\mathbf{x} \frac{\partial}{\partial x_i^\alpha} \left[e^{\beta U(\{\mathbf{x}^\gamma\})} \right] \prod_{\eta} \delta(\mathbf{x}^\eta \cdot \mathbf{p} - w), \quad (\text{B.16})$$

which finally gives

$$\Gamma_\alpha = \frac{\partial}{\partial w^\alpha} \log Z_N = \frac{\beta}{p_i} \left\langle \frac{\partial U}{\partial x_i^\alpha} \right\rangle \quad \forall i. \quad (\text{B.17})$$

The other terms in the thermodynamic Slutsky matrix expression can be simplified in similar ways. Indeed starting from

$$\frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle = \frac{1}{Z_N} \int \mathcal{D}\mathbf{x} x_i^\alpha x_j^\gamma e^{\beta U(\{\mathbf{x}^\eta\})} \frac{\partial}{\partial w^\gamma} \prod_{\eta} \delta(\mathbf{x}^\eta \cdot \mathbf{p} - w^\eta) - \Gamma_\gamma \langle x_i^\alpha x_j^\gamma \rangle, \quad (\text{B.18})$$

and reusing Eq. (B.15) and subsequent iterations by parts, it is straightforward to show that

$$\frac{\partial}{\partial w^\gamma} \langle x_i^\alpha x_j^\gamma \rangle = \frac{1}{p_k} \left[\delta_{ik} \delta_{\alpha\gamma} \langle x_j^\gamma \rangle + \delta_{jk} \langle x_i^\alpha \rangle + \beta \left(\left\langle x_i^\alpha x_j^\gamma \frac{\partial U}{\partial x_k^\gamma} \right\rangle - \langle x_i^\alpha x_j^\gamma \rangle \left\langle \frac{\partial U}{\partial x_k^\gamma} \right\rangle \right) \right] \quad \forall k, \quad (\text{B.19})$$

from which we can also derive

$$\frac{\partial}{\partial w^\gamma} \langle x_i^\alpha \rangle = \frac{1}{p_j} \left[\delta_{ij} \delta_{\alpha\gamma} + \beta \left(\left\langle x_i^\alpha \frac{\partial U}{\partial x_j^\gamma} \right\rangle - \langle x_i^\alpha \rangle \left\langle \frac{\partial U}{\partial x_j^\gamma} \right\rangle \right) \right] \quad \forall j. \quad (\text{B.20})$$

Appendix B. Slutsky matrices

B.1.4 Aggregate Slutsky matrix

We now take

$$\mathcal{S}_{ij} := \frac{\partial}{\partial p_j} \langle \bar{x}_i \rangle + \langle \bar{x}_j \rangle \frac{\partial}{\partial \bar{w}} \langle \bar{x}_i \rangle, \quad (\text{B.21})$$

with $\bar{w}$ the arithmetic mean over the agents' budgets. Now, the above can simply be rewritten as

$$\mathcal{S}_{ij} = \frac{1}{N} \sum_{\alpha} \frac{\partial}{\partial p_j} \langle x_i^{\alpha} \rangle + \frac{1}{N^2} \sum_{\alpha, \gamma} \langle x_j^{\alpha} \rangle \frac{\partial}{\partial \bar{w}} \langle x_i^{\gamma} \rangle. \quad (\text{B.22})$$

The fluctuation-dissipation relation for the first term still holds as before, while for the derivative with respect to the mean budget we have, in general,

$$\frac{\partial}{\partial \bar{w}} \langle x_i^{\gamma} \rangle = \sum_{\eta} \kappa^{\eta} \frac{\partial}{\partial w^{\eta}} \langle x_i^{\gamma} \rangle, \quad (\text{B.23})$$

with

$$\kappa^{\eta} := \frac{\partial w^{\eta}}{\partial \bar{w}}.$$

As a result, we may write

$$\frac{1}{N^2} \sum_{\alpha, \gamma} \langle x_j^{\alpha} \rangle \frac{\partial}{\partial \bar{w}} \langle x_i^{\gamma} \rangle = \frac{1}{N^2} \sum_{\alpha, \gamma, \eta} \kappa^{\eta} \langle x_j^{\alpha} \rangle \frac{\partial}{\partial w^{\eta}} \langle x_i^{\gamma} \rangle = \frac{1}{N} \sum_{\gamma, \eta} \kappa^{\eta} \bar{x}_j \frac{\partial}{\partial w^{\eta}} \langle x_i^{\gamma} \rangle. \quad (\text{B.24})$$

Bringing both expressions together and changing indices appropriately, we find

$$\mathcal{S}_{ij} = -\frac{1}{N} \sum_{\alpha, \gamma} \left[\Gamma_{\gamma} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c + \frac{\partial}{\partial w^{\gamma}} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c + \left(\langle x_j^{\gamma} \rangle - \kappa^{\gamma} \bar{x}_j \right) \frac{\partial}{\partial w^{\gamma}} \langle x_i^{\alpha} \rangle \right]. \quad (\text{B.25})$$

In the limit $\beta \rightarrow \infty$, we have shown that the second term in the sum vanishes. In this case, we therefore recover a symmetric Slutsky matrix whenever $\langle x_j^{\gamma} \rangle = \kappa^{\gamma} \bar{x}_j$. This is for example the case when the budgets are proportional to the mean i.e. $w^{\gamma} = \kappa^{\gamma} \bar{w}$ and all agents have identical preferences.

Now, in a simple model where $w^{\gamma} = \kappa^{\gamma} \bar{w}$, $\forall \gamma$, and all agents have identical preferences, then

$$\langle x_j^{\gamma} \rangle \equiv \kappa^{\gamma} \bar{x}_j$$

and the second contribution to $\mathcal{S}_{ij}$ identically vanishes, so

$$\mathcal{S}_{ij} = -\frac{1}{N} \sum_{\alpha, \gamma} \left[\Gamma_{\gamma} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c + \frac{\partial}{\partial w^{\gamma}} \langle x_i^{\alpha} x_j^{\gamma} \rangle_c \right], \quad (\text{B.26})$$

which is always symmetric when all agents are identical.

Interesting things may happen when agents are more heterogeneous, or in a model where a change of average wealth does not affect all agents in an affine way. Imagine agents are characterized by a label $\kappa^\alpha \in \mathbb{R}$ such that

$$w^\alpha = w_0 F\left(\frac{\kappa^\alpha w}{w_0}\right), \quad (\text{B.27})$$

where w_0 is a certain fixed wealth scale, w a varying parameter and F a possibly non-linear function. We must impose that

$$\frac{\bar{w}}{w_0} = \frac{1}{N} \sum_{\alpha} F\left(\frac{\kappa^\alpha w}{w_0}\right).$$

For example when $F(x) = x$ one finds the above linear model: $w^\alpha = \kappa^\alpha w$ with

$$w = \frac{\bar{w}}{\bar{\kappa}}.$$

More generally, one has

$$\frac{\partial \bar{w}}{\partial w} = \overline{\kappa F'\left(\frac{\kappa w}{w_0}\right)}$$

and

$$\frac{\partial w^\alpha}{\partial w} = \kappa^\alpha F'\left(\frac{\kappa^\alpha w}{w_0}\right),$$

so in general

$$\frac{\partial w^\alpha}{\partial \bar{w}} \neq \frac{w^\alpha}{\bar{w}}$$

and the resulting Slutsky matrix has no longer obvious reasons to be symmetric.

B.1.5 Gaussian fluctuations

We take $\mathbf{x}^\alpha - \langle \mathbf{x}^\alpha \rangle = \delta \mathbf{x}^\alpha$ and set out to calculate $\langle x_i^\alpha x_j^\gamma \rangle_c = \langle \delta x_i^\alpha \delta x_j^\gamma \rangle$. At this stage, for simplicity, we combine the $N \times M$ degrees of freedom in the single column vector $\mathbf{v} = [x_1^1, \dots, x_M^1, \dots, x_1^N, \dots, x_M^N]^\top$ such that the correlations of interest are given by

$$\langle \delta v_{k_1} \delta v_{k_2} \rangle = \frac{e^{\beta U^*}}{Z_N} \int_{-\infty}^{\infty} \left(\prod_{\alpha} \frac{d\hat{\mu}^\alpha}{2\pi} \right) \int_{-\infty}^{\infty} \left(\prod_{k=1}^{M \times N} d\delta v_k \right) \delta v_{k_1} \delta v_{k_2} e^{\frac{\beta}{2} \delta \mathbf{v}^\top \mathbf{H} \delta \mathbf{v} + i \delta \mathbf{v} \cdot \tilde{\mathbf{p}}}, \quad (\text{B.28})$$

with the resized price vector $\tilde{\mathbf{p}} \in \mathbb{R}^{M \times N}$, $\tilde{p}_k = \hat{\mu}^\alpha p_i$, $k = M(\alpha - 1) + i$, and where the product of budget constraints has been rewritten using the Fourier representation of the Dirac δ . In this Gaussian approximation, the partition function can

Appendix B. Slutsky matrices

be calculated exactly with two consecutive integration by parts,

$$\begin{aligned} Z_N &= e^{\beta U^*} \int_{-\infty}^{\infty} \left(\prod_{\alpha} \frac{d\hat{\mu}^{\alpha}}{2\pi} \right) \int_{-\infty}^{\infty} \left(\prod_{k=1}^{M \times N} d\delta v_k \right) e^{\frac{\beta}{2} \delta \mathbf{v}^{\top} \mathbf{H} \delta \mathbf{v} + i \delta \mathbf{v} \cdot \tilde{\mathbf{p}}} \\ &= e^{\beta U^*} \sqrt{\frac{(2\pi)^{M \times N + N}}{\det(-\beta \mathbf{H}) \det(-\mathbf{G}/\beta)}}, \end{aligned} \quad (\text{B.29})$$

with the $N \times N$ matrix $\mathbf{G}$ defined as

$$G^{\alpha\gamma} = \sum_{i,j} p_i (\mathbf{H}^{-1})_{ij}^{\alpha\gamma} p_j. \quad (\text{B.30})$$

Going back to Eqs. (B.28), completing the square in the exponent gives the change of variable $\delta \mathbf{u} = \delta \mathbf{v} + \frac{i}{\beta} \mathbf{H}^{-1} \tilde{\mathbf{p}}$. The integral then reads

$$\begin{aligned} \langle \delta v_{k_1} \delta v_{k_2} \rangle &= \frac{e^{\beta U^*}}{Z_N} \int_{-\infty}^{\infty} \left(\prod_{\alpha} \frac{d\hat{\mu}^{\alpha}}{2\pi} \right) e^{\frac{1}{2\beta} \tilde{\mathbf{p}}^{\top} \mathbf{H}^{-1} \tilde{\mathbf{p}}} \int_{-\infty}^{\infty} \left(\prod_{k=1}^{M \times N} d\delta u_k \right) \\ &\quad \times \left(\delta u_{k_1} \delta u_{k_2} - \frac{1}{\beta^2} (\mathbf{H}^{-1} \tilde{\mathbf{p}})_{k_1} (\mathbf{H}^{-1} \tilde{\mathbf{p}})_{k_2} \right) e^{\frac{\beta}{2} \delta \mathbf{u}^{\top} \mathbf{H} \delta \mathbf{u}}. \end{aligned} \quad (\text{B.31})$$

The first term is straightforward, and simply gives $-\frac{1}{\beta} (\mathbf{H}^{-1})_{k_1 k_2}$ as without the constraint. The second term requires more care, as we have

$$(\mathbf{H}^{-1} \tilde{\mathbf{p}})_{k_1} = \sum_{\ell_1} (\mathbf{H}^{-1})_{k_1 \ell_1} \tilde{p}_{\ell_1} = \sum_{\ell_1} (\mathbf{H}^{-1})_{k_1 \ell_1} \hat{\mu}^{\alpha(\ell_1)} p_{i(\ell_1)}. \quad (\text{B.32})$$

As a result,

$$\begin{aligned} &\int_{-\infty}^{\infty} \left(\prod_{\alpha} \frac{d\hat{\mu}^{\alpha}}{2\pi} \right) (\mathbf{H}^{-1} \tilde{\mathbf{p}})_{k_1} (\mathbf{H}^{-1} \tilde{\mathbf{p}})_{k_2} e^{\frac{1}{2\beta} \hat{\mu}^{\top} \mathbf{G} \hat{\mu}} \\ &= -\beta \sqrt{\frac{(2\pi)^N}{\det(-\mathbf{G}/\beta)}} \sum_{\ell_1, \ell_2} (\mathbf{H}^{-1})_{k_1 \ell_1} p_{i(\ell_1)} (\mathbf{G}^{-1})_{\alpha(\ell_1) \alpha(\ell_2)} (\mathbf{H}^{-1})_{k_2 \ell_2} p_{i(\ell_2)}. \end{aligned} \quad (\text{B.33})$$

Bringing everything together, we finally have

$$\langle \delta v_{k_1} \delta v_{k_2} \rangle = -\frac{1}{\beta} (\mathbf{H}^{-1})_{k_1 k_2} + \frac{1}{\beta} \sum_{\ell_1, \ell_2} (\mathbf{H}^{-1})_{k_1 \ell_1} p_{i(\ell_1)} (\mathbf{G}^{-1})_{\alpha(\ell_1) \alpha(\ell_2)} (\mathbf{H}^{-1})_{k_2 \ell_2} p_{i(\ell_2)}, \quad (\text{B.34})$$

which, after replacing with the correct indices, is the expression given in the main text.

B.2 Numerical methods

B.2.1 General idea

To measure equilibrium properties of the system, we choose to follow a Monte Carlo approach, although a stochastic gradient descent of the global utility (i.e. Langevin dynamics) should be equivalent. The system is first initialized by randomly allocating the N baskets following a uniform distribution. We then employ a Metropolis-Hastings algorithm (see next subsection), as a modification $\Delta \mathbf{x}$ to a randomly selected agent's basket is proposed and accepted with a rate verifying detailed balance [310]. To satisfy the non-negativity of the baskets and to correctly sample the small x_i^α region if high concentration occurs, the move is actually constructed in terms of logarithms. For a randomly selected agent γ , we take $\log(\mathbf{x}^\gamma + \Delta \mathbf{x}) = \log \mathbf{x}^\gamma + \boldsymbol{\xi}$, with $\xi_i \sim \mathcal{N}(0, 1)$, resulting in a log-normally distributed multiplicative noise. The budget constraint is then enforced by a simple rescaling of the resulting basket of goods. Here, it is essential to note that the proposal distribution is not symmetric. As a result, the asymmetry must be accounted for in the acceptance rate, giving in our specific case

$$\mathbb{P}(\mathbf{x}^\gamma \rightarrow \mathbf{x}^\gamma + \Delta \mathbf{x}) = \min \left(1, e^{\beta \Delta U + \sum_i [\log(x_i^\gamma + \delta x_i) - \log x_i^\gamma]} \right),$$

with ΔU the change in the *global* utility⁴¹ caused by the move (see next subsection). From each run, average quantities such as $\langle \mathbf{x}^\alpha \rangle$ are measured by taking arithmetic means in algorithmic time once the system is equilibrated.

B.2.2 Metropolis-Hastings acceptance rate

In its most general form, the Metropolis-Hastings acceptance rate is given by

$$W(\mathbf{x} \rightarrow \mathbf{y}) = \min \left(1, \frac{q(\mathbf{y} | \mathbf{x}) \pi(\mathbf{y})}{q(\mathbf{x} | \mathbf{y}) \pi(\mathbf{x})} \right), \quad (\text{B.35})$$

where $\pi(\mathbf{x})$ is the probability density function we wish to sample, and $q(\mathbf{y} | \mathbf{x})$ is the conditional probability of proposing the state $\mathbf{y}$ given the current state $\mathbf{x}$. In our case, and as usual in statistical mechanics (Sec. 2.1.2), we simply have $\pi(\mathbf{y})/\pi(\mathbf{x}) = e^{\beta \Delta U}$. The ratio of conditional probabilities requires more care. Given the noise $y_i = x_i e^\xi$, with ξ a Gaussian noise, the probability density of which will be noted ρ , we have

$$q(x_i | y_i) = \int_{\mathbb{R}} d\xi \rho(\xi) \delta(y_i - x_i e^\xi) = \frac{1}{y_i} \rho \left(\log \frac{y_i}{x_i} \right), \quad (\text{B.36})$$

⁴¹See Section 4.5.1 for a complete discussion on the consequences of optimizing the global utility.

Appendix B. Slutsky matrices

and very similarly

$$q(y_i | x_i) = \frac{1}{x_i} \rho\left(-\log \frac{y_i}{x_i}\right). \quad (\text{B.37})$$

Given the symmetry of ρ for zero-mean noise as is the case here, and the fact that all dimensions are statistically independent,

$$\frac{q(\mathbf{y} | \mathbf{x})}{q(\mathbf{x} | \mathbf{y})} = \prod_{i=1}^M \frac{q(y_i | x_i)}{q(x_i | y_i)} = e^{\sum_i \log\left(\frac{y_i}{x_i}\right)}, \quad (\text{B.38})$$

yielding the expression in the previous subsection.

B.2.3 Computing the Slutsky matrix

To compute the entries of the Slutsky matrix, one can use the fluctuation-response relations derived here, which do not require derivatives. However, out-of-equilibrium effects can become significant, in which case these relations do not hold. For this reason, it is essential to compute as a check the original expression given in Eq. (4.13). This is not entirely straightforward, as taking discrete derivatives after time averaging may induce some biases. Instead, we rely on so-called *pathwise derivative estimates*, often used in mathematical finance to compute risk sensitivities [311, 312]. In a nutshell, we generate perturbed trajectories based on the original simulation that will use the same random numbers (both for proposed moves and acceptance), and measure finite differences at every step.

B.3 Interacting model

We consider the global utility

$$U(\{\mathbf{x}^\alpha\}) = \sum_{\alpha=1}^N \sum_{i=1}^M a_i \log x_i^\alpha [1 + c(\bar{x}_i)^k], \quad \bar{x}_i = \frac{1}{N} \sum_{\alpha=1}^N x_i^\alpha. \quad (\text{B.39})$$

B.3.1 Non-interacting limit – canonical ensemble

We start by solving the much simplified non-interacting limit, setting $c = 0$. When there are no interactions in general, the canonical partition function entirely decouples and can be written as a product,

$$Z_N = \prod_{\alpha=1}^N z^\alpha, \quad \text{with} \quad z^\alpha = \int_0^\infty \left(\prod_{i=1}^M dx_i^\alpha \right) e^{\beta u(\mathbf{x}^\alpha)} \delta(\mathbf{p} \cdot \mathbf{x}^\alpha - w^\alpha). \quad (\text{B.40})$$

Appendix B. Slutsky matrices

Now, plugging in the desired utility function and rewriting the Dirac delta with its integral representation, we have

$$z^\alpha = \int_0^\infty \left(\prod_{i=1}^M dx_i^\alpha \right) \int_{-i\infty}^{i\infty} \frac{d\mu}{2\pi i} e^{\beta \sum_i a_i \log x_i^\alpha - \mu \sum_i p_i x_i^\alpha + \mu w^\alpha}. \quad (\text{B.41})$$

This integral now decouples in a product over i , and performing the integral on x_i^α with the appropriate change of variable,

$$\begin{aligned} z^\alpha &= \int_{-i\infty}^{i\infty} \frac{d\mu}{2\pi i} e^{\mu w^\alpha} \prod_i \frac{\Gamma(1 + \beta a_i)}{(\mu p_i)^{1+\beta a_i}} \\ &= \left(\prod_i \frac{\Gamma(1 + \beta a_i)}{p_i^{1+\beta a_i}} \right) \mathcal{L}_\mu^{-1} \left\{ \mu^{-\sum_i (1+\beta a_i)} \right\} (w^\alpha) \\ &= \frac{(w^\alpha)^{\sum_i (1+\beta a_i) - 1}}{\Gamma(M + \beta \sum_i a_i)} \prod_{i=1}^M \frac{\Gamma(1 + \beta a_i)}{p_i^{1+\beta a_i}}, \end{aligned} \quad (\text{B.42})$$

where we have used the known inverse Laplace transform of a power law. Taking the product over all agents then gives the result in the text for the partition function.

To find the average $\langle x_i^\alpha \rangle$, we proceed directly, by computing

$$\langle x_i^\alpha \rangle = \frac{1}{Z_N} \left[\int_0^\infty \left(\prod_{i=1}^M dx_i^\alpha \right) x_i^\alpha e^{\beta u(\mathbf{x}^\alpha)} \delta(\mathbf{p} \cdot \mathbf{x}^\alpha - w^\alpha) \right] \prod_{\gamma \neq \alpha} z^\gamma \quad (\text{B.43})$$

i.e. exactly as before except for the agent and product considered, for which we must now calculate

$$\begin{aligned} \int_0^\infty dx_i^\alpha \int_{-i\infty}^{i\infty} \frac{d\mu}{2\pi i} x_i^\alpha e^{\beta a_i \log x_i^\alpha - \mu p_i x_i^\alpha + \mu w^\alpha} &= \int_{-i\infty}^{i\infty} \frac{d\mu}{2\pi i} \frac{\Gamma(2 + \beta a_i)}{(\mu p_i)^{2+\beta a_i}} \\ &= \int_{-i\infty}^{i\infty} \frac{d\mu}{2\pi i} \frac{1 + \beta a_i}{\mu p_i} \frac{\Gamma(1 + \beta a_i)}{(\mu p_i)^{1+\beta a_i}}, \end{aligned} \quad (\text{B.44})$$

where we have used the property $\Gamma(x + 1) = x\Gamma(x)$. The integration over μ is then identical with only the change in the exponent. At this stage, we can plug back everything together, all $\gamma \neq \alpha$, $j \neq i$ terms will cancel out from the Z_N at the denominator, and we are left with

$$\langle x_i^\alpha \rangle = \frac{w^\alpha}{p_i} \frac{1 + \beta a_i}{\sum_k (1 + \beta a_k)}. \quad (\text{B.45})$$

Appendix B. Slutsky matrices

Thanks to this explicit form, the agent-specific Slutsky matrix entries read

$$S_{ij}^\alpha = \frac{w^\alpha}{p_i p_j} \frac{1 + \beta a_i}{\sum_k (1 + \beta a_k)} \left[\frac{1 + \beta a_j}{\sum_k (1 + \beta a_k)} - \delta_{ij} \right], \quad (\text{B.46})$$

which, as predicted by thermodynamic relations, is clearly symmetric, and where the heterogeneous budgets only act as a prefactor altering the magnitude, such that we can write $S_{ij}^\alpha = w^\alpha K_{ij}$. The distribution of S_{ij}^α is then simply a rescaling of the original distribution of w^α in the large N limit.

B.3.2 Finite interactions – grand-canonical ensemble

To tackle finite interactions, we need to place ourselves in the grand-canonical ensemble, where calculations are possible. Enforcing the value of $\bar{x}_i$ to be the arithmetic mean over all agents with a Dirac distribution expressed in integral form, we have

$$\begin{aligned} \mathcal{Z}_N = \int_0^\infty \left(\prod_{\alpha=1}^N \prod_{i=1}^M dx_i^\alpha \right) \int_0^\infty \left(\prod_{i=1}^M d\bar{x}_i \right) \int_{-i\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) \exp \left(\beta \sum_{\alpha,i} a_i \log x_i^\alpha [1 + c(\bar{x}_i)^k] \right. \\ \left. - \beta \mu \sum_{\alpha,i} x_i^\alpha p_i - N \sum_i \bar{x}_i \lambda_i + \sum_{\alpha,i} x_i^\alpha \lambda_i \right). \end{aligned} \quad (\text{B.47})$$

We may begin by integrating over x_i^α . Given the positive support and logarithmic utility, the result is simply a Gamma function for the N identical agents,

$$\mathcal{Z}_N = \int_0^\infty \left(\prod_{i=1}^M d\bar{x}_i \right) \int_{-i\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) \left(\prod_{i=1}^M \frac{\Gamma(1 + \beta a_i [1 + c(\bar{x}_i)^k])}{(\beta \mu p_i - \lambda_i)^{1 + \beta a_i [1 + c(\bar{x}_i)^k]}} \right)^N e^{-N \sum_i \bar{x}_i \lambda_i}. \quad (\text{B.48})$$

The entire integrand can thus be rewritten as an exponential with N as a prefactor,

$$\begin{aligned} \mathcal{Z}_N = \int_0^\infty \left(\prod_{i=1}^M d\bar{x}_i \right) \int_{-i\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) \exp \left(-N \sum_i \left[\bar{x}_i \lambda_i - \log \Gamma(1 + \beta a_i [1 + c(\bar{x}_i)^k]) \right. \right. \\ \left. \left. + (1 + \beta a_i [1 + c(\bar{x}_i)^k]) \log(\beta \mu p_i - \lambda_i) \right] \right), \end{aligned} \quad (\text{B.49})$$

a form which naturally points to the saddle point approximation as $N \rightarrow \infty$. Indeed, the integral over λ_i above may first be rewritten as approximated as

$$\int_{-i\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) e^{-N g(\bar{\mathbf{x}}, \boldsymbol{\lambda})} \sim \sqrt{\frac{2\pi}{N |\det \mathbf{H}^*|}} e^{-N g(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*)}, \quad (\text{B.50})$$

where λ^* is a minimum of g , and $\mathbf{H}^*$ is the Hessian with respect to λ evaluated at that point. At the saddle, we have

$$\frac{\partial g}{\partial \lambda_i} = \bar{x}_i - \frac{1 + \beta a_i [1 + c(\bar{x}_i)^k]}{\beta \mu p_i - \lambda_i^*} = 0 \quad \Rightarrow \quad \lambda_i^* = \beta \mu p_i - \frac{1}{\bar{x}_i} (1 + \beta a_i [1 + c(\bar{x}_i)^k]), \quad (\text{B.51})$$

and the Hessian is entirely diagonal and given by

$$H_{ij} = \frac{1 + \beta a_i [1 + c(\bar{x}_i)^k]}{(\beta \mu p_i - \lambda_i^*)^2} \delta_{ij} \quad \Rightarrow \quad |\det \mathbf{H}^*| = \prod_{i=1}^M \frac{(\bar{x}_i)^2}{1 + \beta a_i [1 + c(\bar{x}_i)^k]}. \quad (\text{B.52})$$

As a result, the partition function can finally be expressed in the desired form,

$$\mathcal{Z}_N = \int_0^\infty d\bar{\mathbf{x}} e^{-\beta F(\bar{\mathbf{x}})}, \quad (\text{B.53})$$

with $F(\bar{\mathbf{x}}) = N f(\bar{\mathbf{x}})$, and where from the above equations we have

$$\begin{aligned} \beta f(\bar{\mathbf{x}}) = & \sum_i [\beta \mu p_i \bar{x}_i - (1 + \beta a_i [1 + c(\bar{x}_i)^k]) [1 + \log \bar{x}_i - \log(1 + \beta a_i [1 + c(\bar{x}_i)^k])] \\ & - \log \Gamma(1 + \beta a_i [1 + c(\bar{x}_i)^k])] \\ & + \underbrace{\frac{1}{2} \frac{\log N}{N} + \frac{1}{N} \left[\sum_i \left[\log \bar{x}_i - \frac{1}{2} \log(1 + \beta a_i [1 + c(\bar{x}_i)^k]) \right] - \frac{1}{2} \log 2\pi \right]}_{o(1)}, \end{aligned} \quad (\text{B.54})$$

which is completely decoupled in between degrees of freedom. Keeping only $O(1)$ terms, we find the expression given in the main text that we will use onward.

In the thermodynamic limit, the state of the system can once again be determined through the saddle point approximation, as the system will reach a minimum of f with overwhelming likelihood. This amounts to solving

$$\begin{aligned} \frac{\partial f}{\partial \bar{x}_i} = & -\frac{1}{\bar{x}_i^*} \left(\frac{1}{\beta} - \mu p_i \bar{x}_i^* + a_i \left[1 + c(\bar{x}_i^*)^k (1 + k[\psi(1 + \beta a_i [1 + c(\bar{x}_i^*)^k]) \right. \right. \\ & \left. \left. - \log(1 + \beta a_i [1 + c(\bar{x}_i^*)^k]) + \log \bar{x}_i^*] \right] \right) = 0, \end{aligned} \quad (\text{B.55})$$

with ψ the digamma function. To determine the nature of stationary points, we will also require

$$\begin{aligned} \frac{\partial^2 f}{\partial \bar{x}_i^2} = & \frac{1}{(\bar{x}_i)^2} \left[\frac{(1/\beta + a_i [1 - c(k-1)(\bar{x}_i)^k])^2}{1/\beta + a_i [1 + c(\bar{x}_i)^k]} - k c a_i (\bar{x}_i)^k \left(\beta k c a_i (\bar{x}_i)^k \psi^{(1)}(1 + \beta a_i [1 + c(\bar{x}_i)^k]) \right. \right. \\ & \left. \left. + [k-1][\psi(1 + \beta a_i [1 + c(\bar{x}_i)^k]) - \log(1 + \beta a_i [1 + c(\bar{x}_i)^k]) + \log \bar{x}_i] \right) \right], \end{aligned} \quad (\text{B.56})$$

Appendix B. Slutsky matrices

where $\psi^{(1)}$ is the first polygamma function. Remarkably, the second derivative is not dependent on the chemical potential μ . Importantly, one can evaluate the partition function by steepest descent to show that, consistent with intuition,

$$\langle x_i \rangle = -\frac{1}{N\beta\mu} \frac{\partial}{\partial p_i} \log \mathcal{Z} = \bar{x}_i^*, \quad (\text{B.57})$$

when $N \rightarrow \infty$, straight from Eq. (B.54) as all derivatives other than those with respect to p_i are zero at the saddle by construction.

$c = 0$ solution Taking $c = 0$, we easily recover the unconcentrated solution

$$\bar{x}_i^* = \frac{w}{p_i} \frac{1 + \beta a_i}{\sum_k (1 + \beta a_k)} \quad (\text{B.58})$$

$\beta \rightarrow \infty$ limit In order to better understand the influence of the different terms on the equilibrium solution of the system, we start by studying the zero temperature, or fully rational, limit of the system. Knowing the asymptotics of the digamma function $\psi(z) \sim \log z - \frac{1}{2z}$, we obtain

$$a_i [1 + c(\bar{x}_i^*)^k (1 + k \log \bar{x}_i^*)] = \mu p_i \bar{x}_i^*, \quad \sum_i \bar{x}_i^* p_i = w, \quad (\text{B.59})$$

i.e. a set of $M + 1$ equations for the M variables $\bar{x}_i^*$ and the correct value of μ for the given budget w . To obtain the solution as a function of c , one can numerically solve the system of equations.

Knowing what the uniform solution is, we may then identify where the phase transition occurs, i.e. what the value of c_∞ is, by studying the stability of the free energy at that point, as it should go from being a minimum to a maximum at the transition. Now in this limit, we have

$$\frac{\partial^2 f}{\partial \bar{x}_i^2} = \frac{a_i}{(\bar{x}_i^*)^2} \left[\frac{(1 - c(k-1)(\bar{x}_i^*)^k)^2}{1 + c(\bar{x}_i^*)^k} - kc(\bar{x}_i^*)^k \left(\frac{kc(\bar{x}_i^*)^k}{1 + c(\bar{x}_i^*)^k} + (k-1) \log \bar{x}_i^* \right) \right]. \quad (\text{B.60})$$

Setting to zero at the saddle, and solving for c_∞ finally yields

$$\frac{1}{c_\infty} = \max_{\bar{x}_i^*} (\bar{x}_i^*)^k [2k - 1 + k(k-1) \log \bar{x}_i^*]. \quad (\text{B.61})$$

In the case where all products are equivalent, we recover the expression in the main text.

Finite β In the general case, the equations for $\bar{x}_i^*$ are much harder to manipulate due to the presence of the highly nonlinear logarithm and digamma functions. However, the method outlined to obtain the critical value of c remains valid: the $c = 0$ solution is to be plugged into the second derivative that is set to zero. Solving numerically (Eq. (B.56)) then yields the critical line in (c, β) shown in the main text and in Fig. 4.2.

B.3.3 Slutsky matrix in the $\beta \rightarrow \infty$ limit

We wish to employ the previously devised Gaussian approximation for the Slutsky matrix at $\beta \rightarrow \infty$ for our interacting model. To do so means evaluating the “thermodynamic” expression of the Slutsky matrix that requires computing the cross-agent term $\langle x_j^\gamma \rangle \frac{\partial}{\partial w^\gamma} \langle x_i^\alpha \rangle$. For our specific choice of utility,

$$\frac{\partial U}{\partial x_j^\gamma} = \frac{a_j}{x_j^\gamma} [1 + c(\bar{x}_j)^k] + kca_j(\bar{x}_j)^{k-1} \overline{\log x_j}, \quad (\text{B.62})$$

with overlines indicating arithmetic averages over the N agents. As we are interested in the $N \rightarrow \infty$ behavior, we can first notice that in the vicinity of the maximum

$$\bar{x}_i = \bar{x}_i^* + \frac{1}{N} \sum_{\alpha} \delta x_i^\alpha = \bar{x}_i^* + O\left(\frac{1}{\sqrt{N}}\right) \quad (\text{B.63})$$

by virtue of the central limit theorem. Similarly,

$$\begin{aligned} \overline{\log x_i} &= \frac{1}{N} \sum_{\alpha} \log(\bar{x}_i^* + \delta x_i^\alpha) = \log \bar{x}_i^* + \frac{1}{N} \sum_{\alpha} \left[\frac{\delta x_i^\alpha}{\bar{x}_i^*} - \frac{1}{2} \left(\frac{\delta x_i^\alpha}{\bar{x}_i^*} \right)^2 + o\left(\frac{1}{\beta}\right) \right] \\ &= \log \bar{x}_i^* + O\left(\frac{1}{\sqrt{N}}\right) + O\left(\frac{1}{\beta}\right). \end{aligned} \quad (\text{B.64})$$

To leading order, we therefore have

$$\frac{\partial U}{\partial x_j^\gamma} = \frac{a_j}{\bar{x}_j^* + \delta x_j^\gamma} [1 + c(\bar{x}_j^*)^k] + kca_j(\bar{x}_j^*)^{k-1} \log \bar{x}_j^* + O\left(\frac{1}{\sqrt{N}}\right). \quad (\text{B.65})$$

As a result, developing to the second order in δx_j^γ

$$\left\langle x_i^\alpha \frac{\partial U}{\partial x_j^\gamma} \right\rangle - \langle x_i^\alpha \rangle \left\langle \frac{\partial U}{\partial x_j^\gamma} \right\rangle = -\frac{a_j}{(\bar{x}_j^*)^2} \langle x_i^\alpha x_j^\gamma \rangle_c [1 + c(\bar{x}_j^*)^k] + O\left(\frac{1}{\sqrt{N}}\right), \quad (\text{B.66})$$

so plugging back this expression in the fluctuation-response relation gives

$$\frac{\partial}{\partial w^\gamma} \langle x_i^\alpha \rangle = \frac{1}{p_j} \left(\delta_{ij} \delta_{\alpha\gamma} - \beta \frac{a_j}{(\bar{x}_j^*)^2} \langle x_i^\alpha x_j^\gamma \rangle_c [1 + c(\bar{x}_j^*)^k] \right) \quad \forall j. \quad (\text{B.67})$$

Appendix B. Slutsky matrices

Given the Gaussian fluctuations scale as $\langle x_i^\alpha x_j^\gamma \rangle_c \sim \beta^{-1}$, this term will clearly have a non-negligible contribution to the Slutsky matrix at $\beta \rightarrow \infty$.

Using the fluctuation-response relation to compute the first term in the Slutsky matrix,

$$\begin{aligned}\Gamma_\gamma &= \frac{\beta}{p_\ell} \left\langle \frac{\partial U}{\partial x_\ell^\gamma} \right\rangle \quad \forall \ell \\ &= \beta \frac{a_\ell}{p_\ell \bar{x}_\ell^*} \left[1 + c(\bar{x}_\ell^*)^k (1 + k \log \bar{x}_\ell^*) + O\left(\frac{1}{\beta}\right) \right],\end{aligned}\tag{B.68}$$

we can finally write, choosing $\ell = j$

$$S_{ij}^\alpha = -\beta \sum_\gamma \langle x_i^\alpha x_j^\gamma \rangle_c \frac{a_j}{p_j \bar{x}_j^*} \left(k c(\bar{x}_j^*)^k \log \bar{x}_j^* + \delta_{\alpha\gamma} [1 + c(\bar{x}_j^*)^k] \right). \tag{B.69}$$

There now remains to compute the pairwise correlations, which requires us to invert the Hessian of the utility function. In our specific case

$$\begin{aligned}H_{ij}^{\alpha\gamma} &= \frac{\partial^2 U}{\partial x_i^\alpha \partial x_j^\gamma} \\ &= - \left[\frac{a_i}{(\bar{x}_i^\gamma)^2} [1 + c(\bar{x}_i)^k] \delta_{\alpha\gamma} - \frac{kc}{N} a_i (\bar{x}_i)^{k-1} \left(\frac{1}{x_i^\alpha} + \frac{1}{x_i^\gamma} \right) + \frac{k(k-1)c}{N} a_i \overline{\log x_i} (\bar{x}_i)^{k-2} \right] \delta_{ij}.\end{aligned}\tag{B.70}$$

Due to the homogeneity in between agents, this $M \times N$ matrix has the very simple structure

$$\mathbf{H} = \begin{bmatrix} \mathbf{A} & 0 & \cdots & 0 \\ 0 & \mathbf{A} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A} \end{bmatrix} + \begin{bmatrix} \mathbf{B} & \mathbf{B} & \cdots & \mathbf{B} \\ \mathbf{B} & \mathbf{B} & \cdots & \mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{B} & \mathbf{B} & \cdots & \mathbf{B} \end{bmatrix}, \tag{B.71}$$

with $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{M \times M}$, both having the additional simplification of being symmetric matrices. We take the *ansatz* that the inverse has an identical structure,

$$\mathbf{H}^{-1} = \begin{bmatrix} \mathbf{F} & 0 & \cdots & 0 \\ 0 & \mathbf{F} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{F} \end{bmatrix} + \begin{bmatrix} \mathbf{D} & \mathbf{D} & \cdots & \mathbf{D} \\ \mathbf{D} & \mathbf{D} & \cdots & \mathbf{D} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{D} & \mathbf{D} & \cdots & \mathbf{D} \end{bmatrix}. \tag{B.72}$$

As such, we have

$$\mathbf{H}\mathbf{H}^{-1} = \begin{bmatrix} \mathbf{A}\mathbf{F} & 0 & \cdots & 0 \\ 0 & \mathbf{A}\mathbf{F} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{A}\mathbf{F} \end{bmatrix} + \begin{bmatrix} \mathbf{E} & \mathbf{E} & \cdots & \mathbf{E} \\ \mathbf{E} & \mathbf{E} & \cdots & \mathbf{E} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{E} & \mathbf{E} & \cdots & \mathbf{E} \end{bmatrix}, \quad \mathbf{E} = \mathbf{A}\mathbf{D} + \mathbf{B}\mathbf{F} + N\mathbf{B}\mathbf{D}, \tag{B.73}$$

Appendix B. Slutsky matrices

obtained from the multiplication of these simple block matrices. As a result, setting $\mathbf{H}\mathbf{H}^{-1} = \mathbf{I}$ amounts to

$$\mathbf{F} = \mathbf{A}^{-1}, \quad \mathbf{D} = -(\mathbf{A} + N\mathbf{B})^{-1}\mathbf{B}\mathbf{A}^{-1}. \quad (\text{B.74})$$

Going back to Eq. (B.70) (evaluated at the maximum $\bar{\mathbf{x}}^*$), we have

$$A_{ij} = -a_i \left(\frac{1 + c(\bar{x}_i^*)^k}{(\bar{x}_i^*)^2} + \frac{kc}{N} (\bar{x}_i^*)^{k-2} [2 + (k-1) \log \bar{x}_i^*] \right) \delta_{ij} \quad (\text{B.75})$$

and

$$B_{ij} = \frac{kc}{N} a_i (\bar{x}_i^*)^{k-2} [2 + (k-1) \log \bar{x}_i^*] \delta_{ij}. \quad (\text{B.76})$$

The diagonal matrix can be inverted (almost) exactly in the large N limit, yielding

$$F_{ij} = -\frac{\delta_{ij}}{a_i} \frac{(\bar{x}_i^*)^2}{1 + c(\bar{x}_i^*)^k} \left[1 - \frac{kc}{N} \frac{(\bar{x}_i^*)^k}{1 + c(\bar{x}_i^*)^k} (2 + (k-1) \log \bar{x}_i^*) + O\left(\frac{1}{N^2}\right) \right], \quad (\text{B.77})$$

and similarly at the leading order

$$\begin{aligned} [(\mathbf{A} + N\mathbf{B})^{-1}]_{ij} = & -\frac{\delta_{ij}}{a_i} \frac{(\bar{x}_i^*)^2}{1 + c(\bar{x}_i^*)^k [1 - k(2 + (k-1) \log \bar{x}_i^*)]} \\ & \times \left[1 - \frac{kc}{N} \frac{(\bar{x}_i^*)^k (2 + (k-1) \log \bar{x}_i^*)}{1 + c(\bar{x}_i^*)^k [1 - k(2 + (k-1) \log \bar{x}_i^*)]} \right], \end{aligned} \quad (\text{B.78})$$

finally giving

$$D_{ij} = -\frac{kc}{N} \frac{\delta_{ij}}{a_i} \frac{(\bar{x}_i^*)^{k+2} (2 + (k-1) \log \bar{x}_i^*)}{(1 + c(\bar{x}_i^*)^k) (1 + c(\bar{x}_i^*)^k [1 - k(2 + (k-1) \log \bar{x}_i^*)])} + O\left(\frac{1}{N^2}\right). \quad (\text{B.79})$$

Using the inverse of the Hessian, we can then compute the correlations of interest. We first need

$$G^{\alpha\gamma} = \sum_{i,j} p_i (\mathbf{H}^{-1})_{ij}^{\alpha\gamma} p_j, \quad (\text{B.80})$$

which can clearly be split in diagonal and off-diagonal parts, therefore the matrix has a similar structure to the Hessian, but reduced in dimension. We have $\mathbf{G} = a\mathbf{I} + b\mathbf{E}$ with $\mathbf{I}$ and $\mathbf{E}$ respectively identity matrix and a matrix full of ones, both of dimension $N \times N$. Taking the same ansatz as before but with scalars, we have

$$\mathbf{G}^{-1} = \frac{1}{a} \mathbf{I} + d\mathbf{E}, \quad \text{with} \quad d = -\frac{b}{a(a + Nb)}, \quad a = \mathbf{p}^\top \mathbf{F} \mathbf{p}, \quad b = \mathbf{p}^\top \mathbf{D} \mathbf{p} \quad (\text{B.81})$$

Appendix B. Slutsky matrices

Bringing everything together, we finally find,

$$\begin{aligned}\langle x_i^\alpha x_j^\gamma \rangle_c &= -\frac{1}{\beta}(\delta_{\alpha\gamma} F_{ii} + D_{ii})\delta_{ij} + \frac{p_i p_j}{\beta} [F_{ii} F_{jj} (\delta_{\alpha\gamma} a^{-1} + d) \\ &\quad + (a^{-1} + Nd)(F_{ii} D_{jj} + D_{ii} F_{jj} + N D_{ii} D_{jj})] \\ &= \frac{1}{\beta} \left[\varphi_{ij} \delta_{\alpha\gamma} + \frac{1}{N} \psi_{ij} \right]\end{aligned}\quad (\text{B.82})$$

with

$$\varphi_{ij} = F_{ii} \delta_{ij} - a^{-1} p_i p_j F_{ii} F_{jj} \quad (\text{B.83})$$

and

$$\frac{1}{N} \psi_{ij} = D_{ii} \delta_{ij} - p_i p_j [d F_{ii} F_{jj} + (a^{-1} + Nd)(F_{ii} D_{jj} + D_{ii} F_{jj} + N D_{ii} D_{jj})] \quad (\text{B.84})$$

the non-interacting and interacting parts respectively. Given that $D_{ii} \sim \frac{1}{N}$ and $d \sim \frac{1}{N}$, it is clear that we have $\varphi_{ij} = O(1)$ and $\psi_{ij} = O(1)$, as expected.

Ultimately, in the $N \rightarrow \infty$, $\beta \rightarrow \infty$ limits, the Slutsky matrix is thus given by

$$S_{ij}^\alpha = S_{ij} = -\frac{a_j}{p_j} \left[k c(\varphi_{ij} + \psi_{ij})(\bar{x}_j^*)^{k-1} \log \bar{x}_j^* + [1 + c(\bar{x}_j^*)^k] \varphi_{ij} \right]. \quad (\text{B.85})$$

The evolution of this analytical expression for some $\mathbf{a}, \mathbf{p}$ with c , compared with numerical simulations, is shown in Figure 4.3 of the main text.

B.3.4 Equivalence of ensembles

We wish to compute

$$\sigma^2 = \left\langle \left(\sum_i x_i^\alpha p_i - w \right)^2 \right\rangle = \left\langle \left(\sum_i x_i^\alpha p_i \right)^2 \right\rangle - w^2. \quad (\text{B.86})$$

To this end, we take an agent-specific perturbation to the chemical potential $\mu \rightarrow \mu + \delta\mu^\alpha$, playing the role of an external field in conventional statistical physics. The grand-canonical partition function then reads

$$\mathcal{Z}_N = \int_0^\infty \left(\prod_{\alpha,i} dx_i^\alpha \right) e^{\beta[U(\{\mathbf{x}\}) - \mu \sum_{\alpha,i} x_i^\alpha p_i - \sum_{\alpha,i} \delta\mu^\alpha x_i^\alpha p_i]}, \quad (\text{B.87})$$

such that

$$\frac{1}{\mathcal{Z}_N} \frac{\partial \mathcal{Z}_N}{\partial \delta\mu^\alpha} = -\beta \left\langle \sum_i x_i^\alpha p_i \right\rangle = -\beta w \quad \text{and} \quad \frac{1}{\mathcal{Z}_N} \frac{\partial^2 \mathcal{Z}_N}{\partial \delta\mu^{\alpha^2}} = \beta^2 \left\langle \left(\sum_i x_i^\alpha p_i \right)^2 \right\rangle, \quad (\text{B.88})$$

giving

$$\sigma^2 = \frac{1}{\beta^2} \frac{\partial^2}{\partial \delta \mu^{\alpha^2}} \log \mathcal{Z}_N. \quad (\text{B.89})$$

As a result, assuming $\delta \mu^\alpha$ to be small, we need to calculate the free energy to the second order in this perturbation. Following the previous calculation, we have

$$\mathcal{Z}_N = \int_0^\infty \left(\prod_{i=1}^M d\bar{x}_i \right) \int_{-\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) \prod_{\alpha, i} \frac{\Gamma(1 + \beta a_i [1 + c(\bar{x}_i)^k])}{(\beta \mu p_i + \beta \delta \mu^\alpha p_i - \lambda_i)^{1 + \beta a_i [1 + c(\bar{x}_i)^k]}} e^{-N \sum_i \bar{x}_i \lambda_i}. \quad (\text{B.90})$$

Expanding at the second order in the perturbation, we may write

$$\mathcal{Z}_N = \int_0^\infty \left(\prod_{i=1}^M d\bar{x}_i \right) \int_{-\infty}^{i\infty} \left(\prod_{i=1}^M \frac{d\lambda_i}{2\pi i} \right) e^{-N[g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda}) + g_1(\bar{\mathbf{x}}, \boldsymbol{\lambda}) \overline{\delta \mu} + g_2(\bar{\mathbf{x}}, \boldsymbol{\lambda}) \overline{\delta \mu^2}]}, \quad (\text{B.91})$$

to be evaluated through a saddle point approximation, with $g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda})$ corresponding to the exponent in the unperturbed expression previously studied,

$$g_1(\bar{\mathbf{x}}, \boldsymbol{\lambda}) = \sum_i \frac{p_i}{\mu p_i - \lambda_i / \beta} (1 + \beta a_i [1 + c(\bar{x}_i)^k]) \quad (\text{B.92})$$

and

$$g_2(\bar{\mathbf{x}}, \boldsymbol{\lambda}) = -\frac{1}{2} \sum_i \frac{p_i^2}{(\mu p_i - \lambda_i / \beta)^2} (1 + \beta a_i [1 + c(\bar{x}_i)^k]), \quad (\text{B.93})$$

while the previously introduced overline notation implies an arithmetic average over agents, i.e. $\overline{\delta \mu} = \frac{1}{N} \sum_\alpha \delta \mu^\alpha$. Now before actually using these expressions, one may use the properties at the saddle to identify the necessary terms to evaluate at the leading order. Indeed, at the saddle in λ , which decouples in i so can be treated in one dimension for now, we will have

$$\boldsymbol{\lambda} = \boldsymbol{\lambda}^* + \boldsymbol{\theta}_1 \overline{\delta \mu} + \boldsymbol{\theta}_2 \overline{\delta \mu^2} + O(\delta \mu^3) \quad (\text{B.94})$$

solution to

$$\frac{\partial g_0}{\partial \lambda_i} + \frac{\partial g_1}{\partial \lambda_i} \overline{\delta \mu} + \frac{\partial g_2}{\partial \lambda_i} \overline{\delta \mu^2} = 0, \quad (\text{B.95})$$

and thus where λ^* corresponds to the previously calculated unperturbed solution. Now to evaluate the integrand at the saddle, we may expand all expressions to the second order, yielding

$$\mathcal{Z}_N = \int_0^\infty \left(\prod_i d\bar{x}_i \right) e^{-N\beta f(\bar{\mathbf{x}})}, \quad (\text{B.96})$$

Appendix B. Slutsky matrices

with

$$\beta f(\bar{\mathbf{x}}) = g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*) + \frac{1}{2} \boldsymbol{\theta}_1^\top \mathbf{H}_\lambda(g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*)) \boldsymbol{\theta}_1 (\bar{\delta\mu})^2 + g_1(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*) \bar{\delta\mu} + g_2(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*) \bar{\delta\mu}^2, \quad (\text{B.97})$$

with $\mathbf{H}_\lambda(g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*))$ the Hessian of g_0 with respect to the λ_i evaluated at the saddle at the leading order, which is diagonal in our decoupled problem, i.e.

$$\mathbf{H}_\lambda(g_0(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*)) = \text{diag} \left. \frac{\partial^2 g_0}{\partial \lambda_i^2} \right|_{\bar{\mathbf{x}}, \boldsymbol{\lambda}^*}. \quad (\text{B.98})$$

We may now evaluate the remaining integrals over $\bar{x}_i$ by taking a new saddle, found at

$$\bar{\mathbf{x}} = \bar{\mathbf{x}}^* + \boldsymbol{\kappa}_1 \bar{\delta\mu} + \boldsymbol{\kappa}_2 \bar{\delta\mu}^2 + \boldsymbol{\kappa}_3 (\bar{\delta\mu})^2, \quad (\text{B.99})$$

the solution to the set of equations $\frac{\partial f}{\partial \bar{x}_i} = 0$. Once again expanding all solutions about the saddle, we finally find

$$\begin{aligned} \beta f(\bar{\mathbf{x}}) &= g_0(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*) + \frac{1}{2} \left[\boldsymbol{\theta}_1^\top \mathbf{H}_\lambda(g_0(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*)) \boldsymbol{\theta}_1 + \boldsymbol{\kappa}_1^\top \mathbf{H}_{\bar{x}}(g_0(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*)) \boldsymbol{\kappa}_1 \right] (\bar{\delta\mu})^2 \\ &\quad + g_1(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*) \bar{\delta\mu} + g_2(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*) \bar{\delta\mu}^2. \end{aligned} \quad (\text{B.100})$$

However, at this stage, one may notice that

$$\frac{\partial^2}{\partial \delta\mu^2} (\bar{\delta\mu})^2 = \frac{2}{N^2}, \quad (\text{B.101})$$

whereas

$$\frac{\partial^2}{\partial \delta\mu^2} \bar{\delta\mu}^2 = \frac{2}{N}. \quad (\text{B.102})$$

At the leading order in N , we therefore simply have to evaluate g_2 at the original saddle. The result then reads

$$\sigma^2 = \frac{2}{\beta^2} g_2(\bar{\mathbf{x}}^*, \boldsymbol{\lambda}^*) = \frac{1}{\beta^2} \sum_i \frac{p_i^2}{(\mu p_i - \lambda_i^*/\beta)^2} (1 + \beta a_i [1 + c(\bar{x}_i^*)^k]) + O\left(\frac{1}{N}\right). \quad (\text{B.103})$$

Plugging in the known expression of λ_i^* at the original saddle, we finally obtain at the thermodynamic limit

$$\sigma^2 = \sum_i \frac{(\bar{x}_i^* p_i)^2}{1 + \beta a_i [1 + c(\bar{x}_i^*)^k]}. \quad (\text{B.104})$$

As expected, this quantity vanishes for $\beta \rightarrow \infty$, as well as for $c \rightarrow \infty \forall \beta > 0$. In the $c < c_{\text{crit}}$ region, plugging in the non-condensed solution $\bar{x}_i^* = \frac{w}{p_i} \frac{1 + \beta a_i}{\sum_i (1 + \beta a_i)}$ yields

$$\sigma^2 = \frac{w^2}{\sum_i (1 + \beta a_i)} = O\left(\frac{1}{M}\right). \quad (\text{B.105})$$

B.4 A Hamiltonian utility function

We now take

$$U(\{\mathbf{x}^\alpha\}) = \sum_{i,\alpha} a_i \log x_i^\alpha + \frac{1}{2} \sum_{\substack{i,\alpha,\gamma \\ \gamma \neq \alpha}} J_i^{\alpha\gamma} (x_i^\alpha)^\rho (x_i^\gamma)^\rho, \quad (\text{B.106})$$

with *symmetric* interactions $J_i^{\gamma\alpha} = J_i^{\alpha\gamma}$, $\forall i, \alpha, \gamma$ and $0 < \rho < 1$.

We expect such a Hamiltonian model to be qualitatively similar to the one studied in this work. To check that this is the case, we may start by taking the *mean-field* case $J_i^{\alpha\gamma} = a_i J/N$ and the $\beta \rightarrow \infty$ limit. Introducing the Lagrange multipliers μ^α , the Lagrangian to minimize is

$$\mathcal{L}(\{\mathbf{x}^\alpha\}, \{\boldsymbol{\mu}^\alpha\}) = \sum_{i,\alpha} a_i \log x_i^\alpha + J \sum_{i,\alpha} a_i (x_i^\alpha)^\rho \overline{x_i^\rho} - \sum_{\alpha} \mu^\alpha \left(\sum_i p_i x_i^\alpha - w^\alpha \right), \quad (\text{B.107})$$

where the overline again means an average over all agents. This yields the following $N \times (M+1)$ equations

$$\frac{\partial \mathcal{L}}{\partial x_i^\alpha} = \frac{a_i}{x_i^\alpha} \left[1 + J \rho (x_i^\alpha)^\rho \overline{x_i^\rho} \right] - \mu^\alpha p_i = 0 \quad (\text{B.108})$$

$$\frac{\partial \mathcal{L}}{\partial \mu^\alpha} = \sum_i p_i x_i^\alpha - w^\alpha = 0. \quad (\text{B.109})$$

Now assuming that all agents have identical budgets, the maximum will be given by $x_i^\alpha = x_i^* \forall \alpha$. The solution then satisfies

$$x_i^* = \frac{w}{p_i} \frac{a_i [1 + J \rho (x_i^*)^{2\rho}]}{\sum_j a_j [1 + J \rho (x_j^*)^{2\rho}]}. \quad (\text{B.110})$$

In the non-interacting case $J = 0$ we recover the $c = 0$ solution of the previous model.

Now, a phase transition will occur due to changes in the sign of the Hessian of the utility, which is given by

$$\left. \frac{\partial^2 U}{\partial x_i^2} \right|_{x_i^*} = -\frac{a_i}{(x_i^*)^2} \left[1 - J \rho (2\rho - 1) (x_i^*)^{2\rho} \right]. \quad (\text{B.111})$$

Clearly, assuming $J > 0$, we require $\rho > \frac{1}{2}$ regardless of the value of other parameters for a transition to occur. In that case, concentration will indeed occur for $J > J_\infty$ the critical value for the $\beta \rightarrow \infty$ solution. For finite β , the critical value will be given by $J_{\text{crit}} \geq J_\infty$ just as in our original model.

Appendix B. Slutsky matrices

Appendix C

SK-game

C.1 Static NMFE

We claim that for $\alpha \ll 1$

$$\tilde{m}_j^\alpha(t) \simeq \alpha \sum_{t'=1}^t (1-\alpha)^{t-t'} m_j(t'). \quad (\text{C.1})$$

Indeed, given the assumption of independence in time, i.e.

$$\langle (S_j(t') - m_j(t')) (S_j(t'') - m_j(t'')) \rangle = \delta_{t,t'} (1 - m_j(t)),$$

we have

$$\begin{aligned} \left\langle \left(\tilde{m}_j^\alpha(t) - \alpha \sum_{t' \leq t} (1-\alpha)^{t-t'} m_j(t') \right)^2 \right\rangle &= \alpha^2 \sum_{t' \leq t} (1-\alpha)^{2(t-t')} \left\langle (S_j(t') - m_j(t'))^2 \right\rangle \\ &\quad + \alpha^2 \sum_{t' \leq t} \sum_{t'' \neq t'} (1-\alpha)^{t-t'} (1-\alpha)^{t-t''} \langle (S_j(t') - m_j(t')) (S_j(t'') - m_j(t'')) \rangle \\ &= \alpha^2 \sum_{t' \leq t} (1-\alpha)^{2(t-t')} (1 - (m_j(t'))^2) \\ &\leq \alpha^2 \sum_{t' \leq t} (1-\alpha)^{2(t-t')} = \frac{\alpha}{2-\alpha} \xrightarrow{\alpha \rightarrow 0} 0. \end{aligned} \quad (\text{C.2})$$

We can then make the *ansatz* that the expected decision reaches a fixed point m_j^* after some time. For sufficiently large t and small but finite values of α , we will therefore have

$$\langle \tilde{m}_j^\alpha(t) \rangle \simeq m_j^* \quad (\text{C.3})$$

with fluctuations characterized by

$$\left\langle (\tilde{m}_j^\alpha(t) - m_j^*)^2 \right\rangle = \frac{\alpha}{2} (1 - (m_j^*)^2) + O(\alpha^2). \quad (\text{C.4})$$

C.2 Fixed point complexity with cooperativity

Using the “direct” approach, the average number of fixed points in the $\beta \rightarrow \infty$ limit with $J_0 > 0$ is given by

$$\begin{aligned}\overline{\mathcal{N}}_J &= \sum_{\{S\}} \prod_i \overline{\Theta\left(S_i \sum_j J_{ij} S_j + S_i \frac{J_0}{N} \sum_j S_j\right)} \\ &= \lim_{\delta \rightarrow 0^+} \sum_{\{S\}} \int_{-\infty}^{\infty} \prod_i \frac{d\mu_i}{2\pi(i\mu_i + \delta)} e^{i \frac{J_0}{N} \sum_{ij} \mu_i S_i S_j} \overline{e^{i \sum_{ij} \mu_i S_i J_{ij} S_j}}.\end{aligned}\quad (\text{C.5})$$

Now, we can rewrite $S_i J_{ij} S_j = \tilde{J}_{ij}$, with $\tilde{J}_{ij}$ a zero mean and identical variance Gaussian variable, as the interactions reweighed by the Ising spins are still Gaussian. Therefore,

$$\overline{\mathcal{N}}_J = \lim_{\delta \rightarrow 0^+} \sum_{\{S\}} \int_{-\infty}^{\infty} \prod_i \frac{d\mu_i}{2\pi(i\mu_i + \delta)} e^{i \frac{J_0}{N} \sum_{ij} \mu_i S_i S_j} \overline{e^{i \sum_{ij} \mu_i \tilde{J}_{ij}}}. \quad (\text{C.6})$$

Performing the average on these reweighed bonds,

$$\begin{aligned}\overline{e^{i \sum_{ij} \mu_i \tilde{J}_{ij}}} &= \exp\left(-\frac{1}{2N} \left(1 - \frac{\varepsilon}{2}\right)^2 \sum_{i < j} (\mu_i + \mu_j)^2 - \frac{1}{2N} \left(\frac{\varepsilon}{2}\right)^2 \sum_{i < j} (\mu_i - \mu_j)^2\right) \\ &= \exp\left(-\frac{1}{4N} v(\varepsilon) \sum_{i \neq j} (\mu_i^2 + \mu_j^2) - \frac{1}{2N} (1 - \varepsilon) \sum_{i \neq j} \mu_i \mu_j\right).\end{aligned}\quad (\text{C.7})$$

Plugging this in the previous expression,

$$\overline{\mathcal{N}}_J = \lim_{\delta \rightarrow 0^+} \sum_{\{S\}} \int_{-\infty}^{\infty} \prod_i \frac{d\mu_i}{2\pi(i\mu_i + \delta)} \exp\left(i \frac{J_0}{N} \sum_{ij} \mu_i S_i S_j - \frac{1}{2} v(\varepsilon) \sum_i \mu_i^2 - \frac{J^2}{2N} (1 - \varepsilon) \left(\sum_i \mu_i\right)^2\right). \quad (\text{C.8})$$

Now, the goal is to completely factorize the expression with respect to the indices i . We therefore introduce a Hubbard-Stratonovitch transform for the squared sum,

$$e^{-\frac{1}{2N} (1-\varepsilon) (\sum_i \mu_i)^2} = \sqrt{\frac{N}{2\pi(1-\varepsilon)}} \int_{-\infty}^{\infty} dx e^{-\frac{N}{2(1-\varepsilon)} x^2 - ix \sum_i \mu_i}, \quad (\text{C.9})$$

and fix the magnetization

$$M = \frac{1}{N} \sum_i S_i \quad (\text{C.10})$$

with an auxiliary variable $\hat{M}$, such that the original expression becomes

$$\begin{aligned} \overline{\mathcal{N}}_J \propto \lim_{\delta \rightarrow 0^+} \sum_{\{S\}} \int_{-\infty}^{\infty} dx \int_0^1 dM \int_{-\infty}^{\infty} \frac{d\hat{M}}{2\pi} \int_{-\infty}^{\infty} \prod_i \frac{d\mu_i}{2\pi(i\mu_i + \delta)} \exp \left(iJ_0 M \sum_i \mu_i S_i \right. \\ \left. - \frac{1}{2} v(\varepsilon) \sum_i \mu_i^2 - ix \sum_i \mu_i + i\hat{M} \sum_i S_i - \frac{N}{2(1-\varepsilon)} x^2 - iN\hat{M}M \right), \end{aligned} \quad (\text{C.11})$$

which is fully decoupled in-between agent indices. As a result, we can finally write

$$\begin{aligned} \overline{\mathcal{N}}_J \propto \int_{-\infty}^{\infty} dx \int_0^1 dM \int_{-\infty}^{\infty} \frac{d\hat{M}}{2\pi} \exp \left(N \underbrace{\left[-\frac{1}{2(1-\varepsilon)} x^2 - iM\hat{M} + \log \mathcal{J} \right]}_{\tilde{\Sigma}} \right) \\ \sim \exp(N\Sigma(\varepsilon, J_0)), \end{aligned} \quad (\text{C.12})$$

with the *annealed* complexity

$$\Sigma(\varepsilon, J_0) = \underset{x, M, \hat{M}}{\text{saddle}} \tilde{\Sigma}(x, M, \hat{M}, \varepsilon, J_0), \quad (\text{C.13})$$

and

$$\begin{aligned} \mathcal{J} &= \sum_{S=\pm 1} e^{i\hat{M}S} \lim_{\delta \rightarrow 0^+} \int_{-\infty}^{\infty} \frac{d\mu}{2\pi(i\mu + \delta)} \exp \left(-\frac{1}{2} v(\varepsilon) \mu^2 + i \left[J_0 M S - x \right] \mu \right) \\ &= \sum_{S=\pm 1} e^{i\hat{M}S} \Phi \left(\frac{J_0 S M - x}{\sqrt{v(\varepsilon)}} \right), \end{aligned} \quad (\text{C.14})$$

where $\Phi(x) = \frac{1}{2} \text{erfc} \left(-\frac{x}{\sqrt{2}} \right)$ is the Gaussian cumulative distribution function. Taking the change of variable $x \rightarrow -\frac{(1-\varepsilon)}{\sqrt{v(\varepsilon)}} x$, we finally obtain

$$\tilde{\Sigma}(x, M, \hat{M}, \varepsilon, J_0) = -\frac{1}{2} \eta x^2 - iM\hat{M} + \log \sum_{S=\pm 1} e^{i\hat{M}S} \Phi \left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} S M \right). \quad (\text{C.15})$$

Appendix C. SK-game

The three saddle equations are then given by

$$\frac{\partial \tilde{\Sigma}}{\partial(i\hat{M})} = -M + \frac{\sum_{S=\pm 1} S e^{i\hat{M}S} \Phi\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)}{\sum_{S=\pm 1} e^{i\hat{M}S} \Phi\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)} = 0, \quad (\text{C.16})$$

$$\frac{\partial \tilde{\Sigma}}{\partial M} = -i\hat{M} + \frac{J_0}{\sqrt{v(\varepsilon)}} \frac{\sum_{S=\pm 1} S e^{i\hat{M}S} \Phi'\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)}{\sum_{S=\pm 1} e^{i\hat{M}S} \Phi\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)} = 0, \quad (\text{C.17})$$

$$\frac{\partial \tilde{\Sigma}}{\partial x} = -\eta x + \eta \frac{\sum_{S=\pm 1} e^{i\hat{M}S} \Phi'\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)}{\sum_{S=\pm 1} e^{i\hat{M}S} \Phi\left(\eta x + \frac{J_0}{\sqrt{v(\varepsilon)}} SM\right)} = 0. \quad (\text{C.18})$$

Taking $J_0 = 0$, we straightforwardly find $M = \hat{M} = 0$ and recover Eq. (6.5)-(6.6). Interestingly, the replica-symmetric solution of Sherrington and Kirkpatrick ($\varepsilon = 0$) yields the self-consistent equation

$$M = \int_{-\infty}^{\infty} \frac{dz}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \text{sign}(J_0 M + z) = 2\Phi(J_0 M) - 1, \quad (\text{C.19})$$

that corresponds to the saddle point $x = \hat{M} = 0$, which is the valid solution only when $J_0 \gg 1$. This replica-symmetric solution predicts a non-zero magnetization to appear for $J_0 = \sqrt{\frac{\pi}{2}} \approx 1.25$ [28], as opposed to $J_0 = 1$ in the correct replica symmetry-breaking calculation [209], see also the Supplementary Material of [313] for a discussion on the topic in the context of maximum likelihood estimators.

Note that in this calculation we have not introduced any stability criterion for the fixed point. In this discrete case, the criterion should be single-flip stability which is not easy to enforce. As a result, the $M = 0$ fixed point is always a valid saddle point to the complexity, even though we expect it to be dynamically unreachable for large enough J_0 . In this annealed calculation we expect that with a correct stability criterion we should lose the stability of the spin-glass saddle point at the replica-symmetric critical value $J_0 = \sqrt{\frac{\pi}{2}}$. As a result, in order to recover what we believe is the correct annealed complexity, we solve the above saddle equations numerically by taking an initial guess far from the $M = 0$ solution.

Doing so for $\varepsilon \leq 0.5$, we find a discontinuous transition towards non-zero magnetization for J_0 slightly above $\sqrt{\frac{\pi}{2}}$ as shown in Fig. C.1, which suggests that this value is indeed compatible with the value we would get with an appropriate stability criterion on the solution. The important result here is that the complexity appears to remain non-zero throughout the mixed phase. Whether or not this




Figure C.1: (a) Magnetization and (b) annealed complexity as a function of the rescaled incentive to cooperate $J_0/\sqrt{v(\varepsilon)}$ for $\varepsilon = \{0, 0.25, 0.5\}$ shown in dark purple, blue and light green respectively. The dotted lines represent the replica-symmetric transition point $\sqrt{\pi/2}$, while the dashed line in (a) represents the replica-symmetric solution for the magnetization. The dot-dashed line is the RSB result close to its critical point $J_0 = 1$, $M \sim (J_0 - 1)^{3/4}$ [209].

annealed approximation is relevant to describe the typical behavior of the system is left for future work.

C.3 Limit cycle complexity with memory

To study the influence of the memory loss rate $\alpha < 1$, we may adapt the method of Hwang *et al.*, although this requires the introduction of either a strong nonlinearity in the exponent and subsequent saddle equations, or of new variables. We opt for the latter, and write the number of cycles of length L as

$$\begin{aligned} \mathcal{N}_L(N, \alpha, \varepsilon) = \sum_{\{S_i(t)\}} \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{t=1}^L dQ_i(t) \right) |\det \mathcal{J}^\alpha|^N \Theta(Q_i(t) S_i(t)) \\ \times \prod_{i,t} \delta \left(Q_i(t+1) - (1-\alpha)Q_i(t) - \alpha \sum_j J_{ij} S_j(t) \right), \end{aligned} \quad (\text{C.20})$$

where the Dirac δ ensures that the dynamics are satisfied at each step, while the second enforces $S_i(t) = \text{sign}(Q_i(t)) \forall t$. The $L \times L$ matrix $\mathcal{J}^\alpha$ is the α -dependent Jacobian ensuring that the zeros of the δ function are correctly weighted, i.e. for $L > 1$

$$\mathcal{J}_{ts}^\alpha = \begin{cases} -(1-\alpha), & \text{if } s = t \pmod{L} \\ 1, & \text{if } s = t+1 \pmod{L} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{C.21})$$

Appendix C. SK-game

The first step is, as usual, to perform the average over the disorder after introducing the integral representation of the Dirac δ ,

$$\begin{aligned} \overline{\mathcal{N}}_L(N, \alpha, \varepsilon) = & |\det \mathcal{J}^\alpha|^N \sum_{\{S_i(t)\}} \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{t=1}^L dQ_i(t) \frac{d\lambda_i(t)}{2\pi} \right) \left(\prod_{i,t} \Theta(S_i(t)Q_i(t)) \right) \\ & \exp \left[-i \sum_{i,t} \lambda_i(t) Q_i(t+1) + i(1-\alpha) \sum_{i,t} \lambda_i(t) Q_i(t) \right. \\ & - \frac{1}{2N} \alpha^2 \left(1 - \frac{\varepsilon}{2}\right)^2 \sum_{i < j} \left(\sum_t [\lambda_i(t) S_j(t) + \lambda_j(t) S_i(t)] \right)^2 \\ & \left. - \frac{1}{2N} \alpha^2 \left(\frac{\varepsilon}{2}\right)^2 \sum_{i < j} \left(\sum_t [\lambda_i(t) S_j(t) - \lambda_j(t) S_i(t)] \right)^2 \right]. \end{aligned} \quad (\text{C.22})$$

The last two terms, resulting from the average on disorder, may be rearranged to give

$$\begin{aligned} -\frac{\alpha^2}{2N} \sum_{t,s} \left(\sum_{i,j} \left[v(\varepsilon) \lambda_i(t) S_j(t) \lambda_i(s) S_j(s) + (1-\varepsilon) \lambda_i(t) S_j(t) \lambda_j(s) S_i(s) \right] \right. \\ \left. - \frac{1}{2} (\varepsilon - 2)^2 \sum_i \lambda_i(t) S_i(t) \lambda_i(s) S_i(s) \right). \end{aligned}$$

Similar to ref. [201], we introduce a set of auxiliary functions,

$$U(t, s) = \frac{1}{N} \sum_i \lambda_i(t) S_i(t) \lambda_i(s) S_i(s), \quad (\text{C.23})$$

$$V(t, s) = \frac{1}{N} \sum_i \lambda_i(t) S_i(s), \quad (\text{C.24})$$

$$R(t, s) = \frac{1}{N} \sum_i \lambda_i(t) \lambda_i(s), \quad (\text{C.25})$$

$$K(t, s) = \frac{1}{N} \sum_i S_i(t) S_i(s), \quad (\text{C.26})$$

such that the last term gives

$$-\frac{1}{2} N \alpha^2 \sum_{t,s} \left[v(\varepsilon) R(t, s) K(t, s) + (1-\varepsilon) V(t, s) V(s, t) - \frac{(2-\varepsilon)^2}{2N} U(t, s) \right].$$

In the limit $N \rightarrow \infty$, the $O(N^{-1})$ term is insignificant and can thus be neglected. The complete expression is then,

$$\begin{aligned}
 \overline{\mathcal{N}}_L(N, \alpha, \varepsilon) = & |\det \mathcal{J}^\alpha|^N \sum_{\{S_i(t)\}} \int_{-\infty}^{\infty} \left(\prod_{i,t} dQ_i(t) \frac{d\lambda_i(t)}{2\pi} \right) \left(\prod_{i,t} \Theta(S_i(t)Q_i(t)) \right) \\
 & \left(\prod_{t,s} dK(t,s) \frac{d\hat{K}(t,s)}{2\pi} dR(t,s) \frac{d\hat{R}(t,s)}{2\pi} dV(t,s) \frac{d\hat{V}(t,s)}{2\pi} \right) \\
 & \exp \left[-i \sum_{i,t} \lambda_i(t) Q_i(t+1) + i(1-\alpha) \sum_{i,t} \lambda_i(t) Q_i(t) \right. \\
 & \quad \left. - \frac{1}{2} N \alpha^2 \sum_{t,s} \left[v(\varepsilon) R(t,s) K(t,s) + (1-\varepsilon) V(t,s) V(s,t) \right] \right. \\
 & \quad \left. - i \sum_{t,s} \hat{K}(t,s) \left(N K(t,s) - \sum_i S_i(t) S_i(s) \right) \right. \\
 & \quad \left. - i \sum_{t,s} \hat{R}(t,s) \left(N R(t,s) - \sum_i \lambda_i(t) \lambda_i(s) \right) \right. \\
 & \quad \left. - i \sum_{t,s} \hat{V}(t,s) \left(N V(t,s) - \sum_i \lambda_i(t) S_i(s) \right) \right].
 \end{aligned} \tag{C.27}$$

Now, one can notice that

$$\int_{-\infty}^{\infty} \left(\prod_{t,s} \frac{dR(t,s)}{2\pi} \right) e^{-\sum_{t,s} i R(t,s) [N \hat{R}(t,s) - \frac{1}{2} i N \alpha^2 v(\varepsilon) K(t,s)]} = \prod_{t,s} \delta \left(N \hat{R}(t,s) - \frac{1}{2} i N \alpha^2 v(\varepsilon) K(t,s) \right), \tag{C.28}$$

which means the integral over $K(t,s)$ now reads,

$$\int_{-\infty}^{\infty} \left(\prod_{t,s} \frac{dR(t,s)}{2\pi} \right) \prod_{t,s} e^{-i N \hat{R}(t,s) R(t,s)} \delta \left(N \hat{R}(t,s) - \frac{1}{2} i N \alpha^2 v(\varepsilon) K(t,s) \right) = e^{-\frac{2N}{\alpha^2 v(\varepsilon)} \sum_{t,s} \hat{R}(t,s) \hat{K}(t,s)} \tag{C.29}$$

For $V(t,s)$, we have to be a bit more careful, as the expression is not linear for all terms, as the diagonal for $t = s$ gives a quadratic term that will have to be treated separately. Noticing that the product $V(t,s)V(s,t)$ is symmetric, we start by considering the $t < s$,

$$\int_{-\infty}^{\infty} \left(\prod_{t < s} \frac{dV(t,s)}{2\pi} \right) e^{-\sum_{t < s} i V(t,s) [N \hat{V}(t,s) - i N \alpha^2 (1-\varepsilon) V(s,t)]} = \prod_{t < s} \delta \left(N \hat{V}(t,s) - i N \alpha^2 (1-\varepsilon) V(s,t) \right). \tag{C.30}$$

Appendix C. SK-game

Now integrating over the $V(t, s)$ for $t > s$,

$$\int_{-\infty}^{\infty} \left(\prod_{t>s} \frac{dV(t, s)}{2\pi} \right) \prod_{t>s} e^{-iN\hat{V}(t, s)V(t, s)} \delta\left(-i\alpha^2(1-\varepsilon)V(s, t) + 2\hat{V}(t, s)\right) = e^{-\frac{N}{\alpha^2(1-\varepsilon)} \sum_{t>s} \hat{V}(t, s)\hat{V}(s, t)}. \quad (\text{C.31})$$

Finally, the diagonal $V(t, t)$ for which the expression is quadratic can be computed with a Gaussian integral,

$$\int_{-\infty}^{\infty} \left(\prod_t dV(t, t) \right) e^{-\frac{1}{2}N\alpha^2(1-\varepsilon) \sum_t V(t, t)^2 - iN \sum_t V(t, t)\hat{V}(t, t)} \sim e^{-\frac{N}{2\alpha^2(1-\varepsilon)} \sum_t \hat{V}(t, t)^2} \quad (\text{C.32})$$

Up to an $O(1)$ multiplicative constant, the complete expression now reads

$$\begin{aligned} \overline{\mathcal{N}}_L(N, \alpha, \varepsilon) \sim & |\det \mathcal{J}^\alpha|^N \sum_{\{S_i(t)\}} \int_{-\infty}^{\infty} \left(\prod_{i,t} dQ_i(t) \frac{d\lambda_i(t)}{2\pi} \right) \left(\prod_{i,t} \Theta(S_i(t)Q_i(t)) \right) \\ & \left(\prod_{t,s} d\hat{K}(t, s) d\hat{R}(t, s) d\hat{V}(t, s) \right) \exp \left(-i \sum_{i,t} \lambda_i(t) Q_i(t+1) \right. \\ & + i(1-\alpha) \sum_{i,t} \lambda_i(t) Q_i(t) + i \sum_{i,t} \nu_i(t) Q_i(t) S_i(t) \\ & - \frac{2N}{\alpha^2 v(\varepsilon)} \sum_{t,s} \hat{R}(t, s) \hat{K}(t, s) - \frac{N}{2\alpha^2(1-\varepsilon)} \sum_{t,s} \hat{V}(t, s) \hat{V}(s, t) \\ & + i \sum_{t,s} \hat{K}(t, s) \sum_i S_i(t) S_i(s) + i \sum_{t,s} \hat{R}(t, s) \sum_i \lambda_i(t) \lambda_i(s) \\ & \left. + i \sum_{t,s} \hat{V}(t, s) \sum_i \lambda_i(t) S_i(s) \right). \end{aligned} \quad (\text{C.33})$$

Taking the changes of variable

$$\begin{aligned} \hat{R}(t, s) &\rightarrow \frac{1}{2}\alpha^2 v(\varepsilon) \hat{R}(t, s) \\ \hat{K}(t, s) &\rightarrow \frac{1}{2}\alpha^2 v(\varepsilon) \hat{K}(t, s) \\ \hat{V}(t, s) &\rightarrow \alpha^2(1-\varepsilon) \hat{V}(t, s) \end{aligned} \quad (\text{C.34})$$

for convenience, we can entirely factorize the problem in N , finally giving

$$\begin{aligned} \overline{\mathcal{N}}_L(N, \alpha, \varepsilon) \sim & \int_{-\infty}^{\infty} \left(\prod_{t,s} d\hat{K}(t, s) d\hat{R}(t, s) d\hat{V}(t, s) \right) \exp \left(N \left[-\frac{1}{2}\alpha^2 v(\varepsilon) \sum_{t,s} \hat{R}(t, s) \hat{K}(t, s) \right. \right. \\ & \left. \left. - \frac{1}{2}\alpha^2(1-\varepsilon) \sum_{t,s} \hat{V}(t, s) \hat{V}(s, t) + \log \mathcal{I}_L + \log |\det \mathcal{J}^\alpha| \right] \right). \end{aligned} \quad (\text{C.35})$$

where

$$\begin{aligned} \mathcal{I}_L = \sum_{\{S(t)\}} \int_{-\infty}^{\infty} \left(\prod_t dQ(t) \frac{d\lambda(t)}{2\pi} \right) \left(\prod_t \Theta(S(t)Q(t)) \right) \exp \left[-i \sum_t \lambda(t)Q(t+1) \right. \\ \left. + i(1-\alpha) \sum_t \lambda(t)Q(t) + \frac{1}{2}\alpha^2 v(\varepsilon) \sum_{t,s} (S(t)S(s)i\hat{K}(t,s) \right. \\ \left. + \lambda(t)\lambda(s)i\hat{R}(t,s)) + \alpha^2(1-\varepsilon) \sum_{t,s} \lambda(t)S(s)i\hat{V}(t,s) \right]. \end{aligned} \quad (\text{C.36})$$

At this stage, we may notice that $S(t)^2 = 1 \forall t$ and as such that the diagonal part of the sum over $K(t,s)$ can be taken out of $\mathcal{I}_L$. Doing so and combining this contribution to the first term of the exponent of Eq. C.35, we have

$$\begin{aligned} \overline{\mathcal{N}}_L(N, \alpha, \varepsilon) \sim \int_{-\infty}^{\infty} \left(\prod_{t,s} d\hat{K}(t,s) d\hat{R}(t,s) d\hat{V}(t,s) \right) \exp \left(N \left[\frac{1}{2}\alpha^2 v(\varepsilon) \sum_t (i\hat{R}(t,t) + 1)i\hat{K}(t,t) \right. \right. \\ \left. \left. + \frac{1}{2}\alpha^2 \left(1 - \varepsilon + \frac{\varepsilon^2}{2} \right) \sum_{s \neq t} i\hat{R}(t,s)i\hat{K}(t,s) \right. \right. \\ \left. \left. - \frac{1}{2}\alpha^2(1-\varepsilon) \sum_{t,s} \hat{V}(t,s)\hat{V}(s,t) + \log \mathcal{I}_L + \log |\det \mathcal{J}^\alpha| \right] \right). \end{aligned} \quad (\text{C.37})$$

Now, the first term in the exponent can be integrated over the $\hat{K}(t,t)$ exactly, yielding a product of δ functions fixing

$$i\hat{R}(t,t) = -1 \quad \forall t, \quad (\text{C.38})$$

as expected and introducing only a sub-dominant correction $O\left(\frac{\log N}{N}\right)$ in the complexity.

In the $N \rightarrow \infty$ limit, the complexity is then finally given by

$$\begin{aligned} \Sigma_L(\alpha, \varepsilon) = \text{saddle}_{\hat{R}, \hat{K}, \hat{V}} \left\{ \frac{1}{2}\alpha^2 v(\varepsilon) \sum_{s \neq t} i\hat{R}(t,s)i\hat{K}(t,s) - \frac{1}{2}\alpha^2(1-\varepsilon) \sum_{t,s} \hat{V}(t,s)\hat{V}(s,t) \right. \\ \left. + \log \mathcal{I}_L + \log |\det \mathcal{J}^\alpha| \right\}, \end{aligned} \quad (\text{C.39})$$

Appendix C. SK-game

with now

$$\begin{aligned} \mathcal{I}_L = \sum_{\{S(t)\}} \int_{-\infty}^{\infty} \left(\prod_t dQ(t) \frac{d\lambda(t)}{2\pi} \right) \prod_t \Theta(S(t)Q(t)) \exp \left[-i \sum_t \lambda(t)Q(t+1) \right. \\ \left. + i(1-\alpha) \sum_t \lambda(t)Q(t) + \alpha^2(1-\varepsilon) \sum_{t,s} \lambda(t)S(s)i\hat{V}(t,s) \right. \\ \left. - \frac{1}{2}\alpha^2 v(\varepsilon) \left[\sum_t \lambda(t)^2 - \sum_{s \neq t} (S(t)S(s)i\hat{K}(t,s) - \lambda(t)\lambda(s)i\hat{R}(t,s)) \right] \right]. \end{aligned} \quad (\text{C.40})$$

This integral may be rewritten as

$$\begin{aligned} \mathcal{I}_L = \sum_{\{S(t)\}} e^{\frac{1}{2}\alpha^2 v(\varepsilon) \sum_{s \neq t} S(t)S(s)i\hat{K}(t,s)} \int_{-\infty}^{\infty} \left(\prod_t \frac{dQ(t)}{\sqrt{2\pi}} \frac{d\lambda(t)}{\sqrt{2\pi}} \right) \prod_t \Theta(S(t)Q(t)) \\ \exp \left(-\frac{1}{2} \boldsymbol{\lambda}^\top \mathbf{A} \boldsymbol{\lambda} - i \mathbf{b}^\top \boldsymbol{\lambda} \right). \end{aligned} \quad (\text{C.41})$$

with the $L \times L$ matrix $\mathbf{A}$ constituted of

$$A(t,s) = \alpha^2 v(\varepsilon) (\delta_{ts} - (1 - \delta_{ts})i\hat{R}(t,s)), \quad (\text{C.42})$$

and $\mathbf{b} \in \mathbb{R}^L$,

$$b(t) = Q(t+1) - (1-\alpha)Q(t) - \alpha^2(1-\varepsilon) \sum_s \hat{V}(t,s)S(s), \quad (\text{C.43})$$

such that we in fact have

$$\mathbf{b} = \mathcal{J}^\alpha \mathbf{Q} + \mathbf{c}, \quad c(t) = -\alpha^2(1-\varepsilon) \sum_s \hat{V}(t,s)S(s). \quad (\text{C.44})$$

Computing the Gaussian integral on $\boldsymbol{\lambda}$,

$$\begin{aligned} \mathcal{I}_L = \sum_{\{S(t)\}} \frac{e^{\frac{1}{2}\alpha^2 v(\varepsilon) \sum_{s \neq t} S(t)S(s)i\hat{K}(t,s)}}{\sqrt{\det \mathbf{A}}} \int_{-\infty}^{\infty} \left(\prod_t \frac{dQ(t)}{\sqrt{2\pi}} \right) \prod_t \Theta(S(t)Q(t)) \\ \exp \left(-\frac{1}{2} (\mathcal{J}^\alpha \mathbf{Q} + \mathbf{c})^\top \mathbf{A}^{-1} (\mathcal{J}^\alpha \mathbf{Q} + \mathbf{c}) \right). \end{aligned} \quad (\text{C.45})$$

Taking the change of variable $\mathbf{u} = \mathcal{J}^\alpha \mathbf{Q} + \mathbf{c}$, the above becomes

$$\begin{aligned} \mathcal{I}_L = \sum_{\{S(t)\}} \frac{e^{\frac{1}{2}\alpha^2 v(\varepsilon) \sum_{s \neq t} S(t)S(s)i\hat{K}(t,s)}}{|\det \mathcal{J}^\alpha| \sqrt{\det \tilde{\mathbf{A}}}} \int_{-\infty}^{\infty} \left(\prod_t \frac{dQ(t)}{\sqrt{2\pi}} \right) \\ \prod_t \Theta(S(t)(\mathbf{Q} - (\mathcal{J}^\alpha)^{-1} \mathbf{c})(t)) \exp \left(-\frac{1}{2} \mathbf{Q}^\top \tilde{\mathbf{A}}^{-1} \mathbf{Q} \right), \end{aligned} \quad (\text{C.46})$$

where $\tilde{\mathbf{A}} = (\mathcal{J}^\alpha)^{-1} \mathbf{A} (\mathcal{J}^\alpha)^{-1}$. As a result, the $|\det \mathcal{J}^\alpha|$ contributions in the complexity cancel out, and we find ourselves with

$$\Sigma_L(\alpha, \varepsilon) = \text{saddle}_{\hat{R}, \hat{K}, \hat{V}} \left\{ \frac{1}{2} \alpha^2 v(\varepsilon) \sum_{s \neq t} i \hat{R}(t, s) i \hat{K}(t, s) - \frac{1}{2} \alpha^2 (1 - \varepsilon) \sum_{t, s} \hat{V}(t, s) \hat{V}(s, t) + \log \mathcal{I}_L \right\}, \quad (\text{C.47})$$

with now

$$\mathcal{I}_L = \sum_{\{S(t)\}} \frac{e^{\frac{1}{2} \alpha^2 v(\varepsilon) \sum_{s \neq t} S(t) S(s) i \hat{K}(t, s)}}{\sqrt{\det \tilde{\mathbf{A}}}} \int_{-\infty}^{\infty} \left(\prod_t \frac{dQ(t)}{\sqrt{2\pi}} \right) \prod_t \Theta(S(t)(\mathbf{Q} - (\mathcal{J}^\alpha)^{-1} \mathbf{c})(t)) \exp \left(-\frac{1}{2} \mathbf{Q}^\top \tilde{\mathbf{A}}^{-1} \mathbf{Q} \right). \quad (\text{C.48})$$

The integral is challenging to study in generality, as the non-diagonal nature of $\mathcal{J}^\alpha$ and $\mathbf{A}$ means that we cannot factorize the integrand. We now move on to specific cases of interest. The equations can be further simplified through the rescalings $\alpha^2 v(\varepsilon) \hat{K}(t, s) \rightarrow \hat{K}(t, s)$, $\alpha \sqrt{v(\varepsilon)} \hat{V}(t, s) \rightarrow \hat{V}(t, s)$ and $Q(t) \rightarrow \alpha \sqrt{v(\varepsilon)}$, in which case we finally get

$$\Sigma_L(\alpha, \eta) = \text{saddle}_{\hat{R}, \hat{K}, \hat{V}} \left\{ \sum_{s < t} i \hat{R}(t, s) i \hat{K}(t, s) - \frac{\eta}{2} \sum_{t, s} \hat{V}(t, s) \hat{V}(s, t) + \log \mathcal{I}_L \right\}, \quad (\text{C.49})$$

$$\mathcal{I}_L = \sum_{\{S(t)\}} e^{\sum_{s < t} S(t) S(s) i \hat{K}(t, s)} \Psi_L(\Gamma_1(\alpha, \eta), \dots, \Gamma_L(\alpha, \eta); \mathbf{C}), \quad (\text{C.50})$$

where $\Psi_L(x_1, \dots, x_L; \mathbf{C})$ is the cumulative distribution of an L -dimensional Gaussian with a zero mean vector and covariance matrix $\mathbf{C}$ evaluated up to x_t , $t \in \{1, \dots, L\}$. In our case, the upper bounds of integration are given by $\mathbf{\Gamma} = \mathbf{S} \circ (\mathcal{J}^\alpha)^{-1} \mathbf{c}$, with now

$$c(t) = \eta \sum_s \hat{V}(t, s) S(s), \quad (\text{C.51})$$

while the covariance matrix must be taken with care as the off-diagonal elements must be adapted to the presence of $S(t)$ in the Heaviside step function. As a result, off-diagonal elements must be multiplied by $S(t)S(s)$. As a result, we have

$$C(t, s) = \begin{cases} \tilde{A}(t, t) & \text{for } t = s \\ S(t)S(s)\tilde{A}(t, s) & \text{for } t \neq s \end{cases} \quad (\text{C.52})$$

with $\tilde{\mathbf{A}} = (\mathcal{J}^\alpha)^{-1} \mathbf{A} (\mathcal{J}^\alpha)^{-1}$ and

$$A(t, s) = \delta_{ts} - (1 - \delta_{ts}) i \hat{R}(t, s). \quad (\text{C.53})$$

Appendix C. SK-game

C.3.1 Fixed points

We can start by checking that we recover the known result for $L = 1$. In this case, $\mathcal{J}^\alpha = \alpha$, and we only have to solve for a scalar $\hat{V}(1, 1) = x$. Here,

$$\mathcal{I}_L = 2\Psi_1\left(\frac{\eta}{\alpha}x; \frac{1}{\alpha^2}\right) = 2\Phi(\eta x), \quad (\text{C.54})$$

therefore the saddle point equation becomes

$$\Sigma_1(\alpha, \eta) := \Sigma_{\text{FP}}(\eta) = \max_x \left\{ -\frac{1}{2}x^2 + \log 2 + \log \Phi(\eta x) \right\}, \quad (\text{C.55})$$

from which the expression in the main text can immediately be recovered.

C.3.2 Two-cycles

There are now six variables to solve for: $i\hat{R}$, $i\hat{K}$, $\hat{V}_1 = \hat{V}(1, 1)$, $\hat{V}_2 = \hat{V}(2, 2)$, $\hat{V}_{12} = \hat{V}(1, 2)$ and $\hat{V}_{21} = \hat{V}(2, 1)$. In this case, the Jacobian matrix is given by

$$\mathcal{J}^\alpha = \begin{bmatrix} -(1-\alpha) & 1 \\ 1 & -(1-\alpha) \end{bmatrix} \Rightarrow (\mathcal{J}^\alpha)^{-1} = \frac{1}{1-(1-\alpha)^2} \begin{bmatrix} 1-\alpha & 1 \\ 1 & 1-\alpha \end{bmatrix}, \quad (\text{C.56})$$

giving the bounds of integration

$$\mathbf{\Gamma} = \frac{\eta}{1-(1-\alpha)^2} \begin{bmatrix} \hat{V}_{21} + S_1 S_2 \hat{V}_2 + (1-\alpha)(\hat{V}_1 + S_1 S_2 \hat{V}_{12}) \\ \hat{V}_{12} + S_1 S_2 \hat{V}_1 + (1-\alpha)(\hat{V}_2 + S_1 S_2 \hat{V}_{21}) \end{bmatrix} \quad (\text{C.57})$$

and the covariance matrix

$$\mathbf{C} = \frac{1}{(1-(1-\alpha)^2)^2} \begin{bmatrix} 1+(1-\alpha)^2-2(1-\alpha)i\hat{R} & S_1 S_2(2(1-\alpha)-(1+(1-\alpha)^2)i\hat{R}) \\ S_1 S_2(2(1-\alpha)-(1+(1-\alpha)^2)i\hat{R}) & 1+(1-\alpha)^2-2(1-\alpha)i\hat{R} \end{bmatrix}. \quad (\text{C.58})$$

Now, it is immediately apparent that $\mathcal{I}_L$ is a function of the product $S_1 S_2 = \pm 1$ rather than S_1 and S_2 . As a result, the complexity is given by

$$\Sigma_2(\alpha, \eta) = \underset{\hat{R}, \hat{K}, \hat{V}_1, \hat{V}_2, \hat{V}_{12}, \hat{V}_{21}}{\text{saddle}} \left\{ i\hat{R}i\hat{K} - \frac{\eta}{2}(\hat{V}_1^2 + \hat{V}_2^2 + 2\hat{V}_{12}\hat{V}_{21}) + \log 2 + \log \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) \right\}. \quad (\text{C.59})$$

The six saddle point equations read

$$i\hat{R} \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) + \sum_{S=\pm 1} S e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) = 0, \quad (\text{C.60})$$

$$\begin{aligned}
 & \left(i\hat{K} + \frac{i\hat{R}}{1 - (i\hat{R})^2} \right) \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) \\
 & - \frac{1}{2} \sum_{S=\pm 1} e^{Si\hat{K}} \int_{-\infty}^{\Gamma_1} dQ_1 \int_{-\infty}^{\Gamma_2} dQ_2 \mathbf{Q}^\top \frac{\partial \mathbf{C}^{-1}}{\partial i\hat{R}} \mathbf{Q} \frac{e^{-\frac{1}{2} \mathbf{Q}^\top \mathbf{C}^{-1} \mathbf{Q}}}{2\pi \sqrt{\det \mathbf{C}}} = 0
 \end{aligned} \tag{C.61}$$

$$\begin{aligned}
 & -\eta \hat{V}_1 \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) + \frac{\eta(1-\alpha)}{1-(1-\alpha)^2} \sum_{S=\pm 1} e^{Si\hat{K}} \int_{-\infty}^{\Gamma_2} dQ_2 \frac{e^{-\frac{1}{2} \mathbf{Q}_2^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_2^*}}{2\pi \sqrt{\det \mathbf{C}}} \\
 & + \frac{\eta}{1-(1-\alpha)^2} \sum_{S=\pm 1} S e^{Si\hat{K}} \int_{-\infty}^{\Gamma_1} dQ_1 \frac{e^{-\frac{1}{2} \mathbf{Q}_1^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_1^*}}{2\pi \sqrt{\det \mathbf{C}}} = 0
 \end{aligned} \tag{C.62}$$

$$\begin{aligned}
 & -\eta \hat{V}_2 \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) + \frac{\eta}{1-(1-\alpha)^2} \sum_{S=\pm 1} S e^{Si\hat{K}} \int_{-\infty}^{\Gamma_2} dQ_2 \frac{e^{-\frac{1}{2} \mathbf{Q}_2^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_2^*}}{2\pi \sqrt{\det \mathbf{C}}} \\
 & + \frac{\eta(1-\alpha)}{1-(1-\alpha)^2} \sum_{S=\pm 1} e^{Si\hat{K}} \int_{-\infty}^{\Gamma_1} dQ_1 \frac{e^{-\frac{1}{2} \mathbf{Q}_1^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_1^*}}{2\pi \sqrt{\det \mathbf{C}}} = 0
 \end{aligned} \tag{C.63}$$

$$\begin{aligned}
 & -\eta \hat{V}_{21} \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) + \frac{\eta(1-\alpha)}{1-(1-\alpha)^2} \sum_{S=\pm 1} S e^{Si\hat{K}} \int_{-\infty}^{\Gamma_2} dQ_2 \frac{e^{-\frac{1}{2} \mathbf{Q}_2^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_2^*}}{2\pi \sqrt{\det \mathbf{C}}} \\
 & + \frac{\eta}{1-(1-\alpha)^2} \sum_{S=\pm 1} e^{Si\hat{K}} \int_{-\infty}^{\Gamma_1} dQ_1 \frac{e^{-\frac{1}{2} \mathbf{Q}_1^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_1^*}}{2\pi \sqrt{\det \mathbf{C}}} = 0
 \end{aligned} \tag{C.64}$$

$$\begin{aligned}
 & -\eta \hat{V}_{12} \sum_{S=\pm 1} e^{Si\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) + \frac{\eta}{1-(1-\alpha)^2} \sum_{S=\pm 1} e^{Si\hat{K}} \int_{-\infty}^{\Gamma_2} dQ_2 \frac{e^{-\frac{1}{2} \mathbf{Q}_2^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_2^*}}{2\pi \sqrt{\det \mathbf{C}}} \\
 & + \frac{\eta(1-\alpha)}{1-(1-\alpha)^2} \sum_{S=\pm 1} S e^{Si\hat{K}} \int_{-\infty}^{\Gamma_1} dQ_1 \frac{e^{-\frac{1}{2} \mathbf{Q}_1^*{}^\top \mathbf{C}^{-1} \mathbf{Q}_1^*}}{2\pi \sqrt{\det \mathbf{C}}} = 0,
 \end{aligned} \tag{C.65}$$

Appendix C. SK-game

with $\mathbf{Q}_1^* = [Q_1 \quad \Gamma_2]^\top$ and similarly for $\mathbf{Q}_2^*$, and

$$\begin{aligned} \frac{\partial \mathbf{C}^{-1}}{\partial i\hat{R}} &= \frac{1}{(1 - (i\hat{R})^2)^2} \\ &\times \begin{bmatrix} 2i\hat{R}(1 + (1 - \alpha)^2) - 2(1 - \alpha)((i\hat{R})^2 + 1) & S(-4i\hat{R}(1 - \alpha) + (1 + (1 - \alpha)^2)((i\hat{R})^2 + 1)) \\ S(-4i\hat{R}(1 - \alpha) + (1 + (1 - \alpha)^2)((i\hat{R})^2 + 1)) & 2i\hat{R}(1 + (1 - \alpha)^2) - 2(1 - \alpha)((i\hat{R})^2 + 1) \end{bmatrix}. \end{aligned} \quad (\text{C.66})$$

In the specific case where $\alpha = 1$, we expect two-cycles with a zero overlap in between consecutive steps, meaning $i\hat{R} = 0$. As a result, the matrix $\mathbf{C}$ is the identity matrix, and thus

$$\Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) = \Phi\left(\eta(\hat{V}_{21} + S\hat{V}_2)\right) \Phi\left(\eta(\hat{V}_{12} + S\hat{V}_1)\right). \quad (\text{C.67})$$

In order for $i\hat{R} = 0$ to satisfy Eq. (C.60), we then require $i\hat{K} = 0$ and $\hat{V}_1 = \hat{V}_2 = 0$ such that $\sum_{S=\pm 1} S e^{i\hat{K}} \Psi_2(\Gamma_1, \Gamma_2; \mathbf{C}) = 0$. Given

$$\left. \frac{\partial \mathbf{C}^{-1}}{\partial i\hat{R}} \right|_{i\hat{R}=0} = \begin{bmatrix} 0 & S \\ S & 0 \end{bmatrix} \quad (\text{C.68})$$

for $\alpha = 1$, the solution $i\hat{K} = 0$ satisfies the saddle point equation (C.61) while the independence of the integrands on S ensures that $\hat{V}_1 = \hat{V}_2 = 0$ are compatible with Eqs. (C.62)-(C.63). Eqs. (C.64)-(C.65) then suggest the ansatz $\hat{V}_{12} = \hat{V}_{21} = x$, in which case the problem is finally reduced to

$$\begin{aligned} \Sigma_2(\alpha = 1, \eta) &= \max_x \{-\eta x^2 + 2 \log 2 + 2 \log \Phi(\eta x)\} \\ &= 2\Sigma_{\text{FP}}(\eta), \end{aligned} \quad (\text{C.69})$$

recovering the known solution of [201, 216].

C.4 Derivation of the DMFT equations

We start from the $N \gg 1$ discrete difference equations to which we have added an external field $h_i(t)$

$$Q_i(t+1) = (1-\alpha)Q_i(t) + \alpha \sum_j J_{ij}m_j(t) + \alpha \eta_i(t) + h_i(t), \quad \langle \eta_i(t) \eta_j^s \rangle \approx v(\varepsilon)(1-q(t))\delta_{t,s}\delta_{i,j} \quad (\text{C.70})$$

and introduce a new agent at index $i = 0$. The influence of this newly introduced agent on the dynamic equation of agents $i \neq 0$ is then

$$\sum_j J_{ij}m_j(t) + h_i(t) \rightarrow \sum_j J_{ij}m_j(t) + J_{i0}m_0(t) + h_i(t) \quad (\text{C.71})$$

Given N is large, we consider the response at the linear order, meaning that the expected decision of all agents $i > 0$ becomes

$$m_i(t) \rightarrow m_i(t) + \sum_{s < t} \sum_j \underbrace{\frac{\partial m_i(t)}{\partial h_j(s)} \Big|_{h=0}}_{\chi_{ij}(t,s)} J_{j0} m_0(s), \quad (\text{C.72})$$

as $J_{j0} m_0(t)$ can be seen as the modification of the effective field “felt” by all agents j , and $\chi_{ij}(t, s)$ is simply the linear response function to a small external field. The dynamics that the newly introduced agent follows is then given by

$$Q_0(t+1) = (1-\alpha)Q_0(t) + \alpha \sum_{i>0} J_{0i} m_i(t) + \sum_{s < t} \left(\sum_{ij} J_{0i} \chi_{ij}(t, s) J_{j0} \right) m_0(s) + \alpha \eta_0(t) + h_0(t). \quad (\text{C.73})$$

The sum on i, j can be split in its diagonal and off-diagonal parts. On the diagonal, we assume the central limit theorem holds, yielding

$$\sum_i J_{0i} \chi_{ii}(t, s) J_{i0} = N \left[\overline{\langle J_{0i} \chi_{ii}(t, s) J_{i0} \rangle} + O\left(\frac{1}{\sqrt{N}}\right) \right] \approx (1 - \varepsilon) \langle \chi_{ii}(t, s) \rangle, \quad (\text{C.74})$$

while the off-diagonal contribution will be sub-dominant as its mean will be zero given non-opposing entries in the interaction matrix are uncorrelated. We can also assume the CLT is valid for the sum on indices i at the leading order, in which case

$$\sum_i J_{0i} m_i(t) \approx 0 + v(\varepsilon) \xi_0(t), \quad (\text{C.75})$$

where ξ_0 is a Gaussian of zero mean and correlated in time as $\langle \xi_0(t) \xi_0(s) \rangle = C_0(t, s) = \langle m_0(t) m_0(s) \rangle$. Bringing everything together, one realizes that there is no cross contributions between agents at the leading order and thus that we can drop the index 0 and recover the equation in the main text

$$Q(t+1) = (1-\alpha)Q(t) + \alpha^2(1-\varepsilon) \sum_{s < t} G(t, s) m(s) + \alpha \phi(t) + \alpha h(t), \quad (\text{C.76})$$

with a new noise term combining the original thermal-like fluctuations and the effective contribution from the disorder averaging, with $\langle \phi(t) \rangle = 0$, and

$$\langle \phi(t) \phi(s) \rangle = v(\varepsilon) [C(t, s) + (1 - q(t)) \delta_{t,s}], \quad (\text{C.77})$$

and where the memory kernel and correlation function are to be determined self-consistently,

$$G(t, s) = \langle \chi_{ii}(t, s) \rangle = \left\langle \frac{\partial m(t)}{\partial h(s)} \Big|_{h=0} \right\rangle, \quad C(t, s) = \langle m(t) m(s) \rangle. \quad (\text{C.78})$$

Appendix C. *SK-game*

Note that in order to eliminate the external field, we have expressed the susceptibility with the noise term, resulting in a rescaling $G(t, s) \rightarrow \alpha G(t, s)$. From this expression, the continuous limit can be taken, changing the sum in time weighted by α into an integral,

$$\frac{\alpha}{2}\ddot{Q}(t) + \dot{Q}(t) = -Q(t) + (1 - \varepsilon) \int_0^t ds G(t, s)m(s) + \phi(t) + h(t). \quad (\text{C.79})$$

Appendix D

Schelling model

This section reproduces the relevant Appendices of [5] and thus includes contributions from R. Zakine and A.-C. Becharat.

D.1 Lyapunov function for non-local moves

In this section, we explicitly show that the non-local mean-field dynamics correspond to the relaxation of a free energy functional when the utility is linear. This result is a direct consequence of the dynamics following the logit rule being in detailed balance with the Gibbs-Boltzmann measure $e^{-\beta\mathcal{F}}$. Starting from the pairwise Hamiltonian

$$\mathcal{H} = -\frac{\nu}{2} \sum_{\mathbf{r}, \mathbf{r}'} n(\mathbf{r}) G_{\sigma}(\mathbf{r} - \mathbf{r}') n(\mathbf{r}'), \quad (\text{D.1})$$

which can be shown to satisfy Eq. (7.5) when $u(\phi) = \nu\phi$, taking the continuous limit and accounting for the entropic contribution, one finds the free energy functional

$$\mathcal{F}[\rho] = -\frac{\nu}{2} \int dx dy \rho(x) G_{\sigma}(x - y) \rho(y) + TS[\rho], \quad (\text{D.2})$$

with the entropy $S[\rho] = \int [\rho \log(\rho) + (1 - \rho) \log(1 - \rho)]$. We define the chemical potential

$$\mu(x) = \frac{\delta \mathcal{F}}{\delta \rho(x)} = -\nu\phi(x) + T \log \left(\frac{\rho(x)}{1 - \rho(x)} \right), \quad (\text{D.3})$$

such that inverting the equation yields

$$\rho(x) = \frac{e^{\frac{\beta}{2}(\mu(x) + \nu\phi(x))}}{D(x)}, \quad (\text{D.4})$$

Appendix D. Schelling model

with $D(x) \equiv 2 \cosh[\frac{\beta}{2}(\mu(x) + \nu\phi(x))]$. Then, computing the total time derivative on the functional yields

$$\begin{aligned} \frac{d\mathcal{F}}{dt} &= \int_x \frac{\delta\mathcal{F}}{\delta\rho(x)} \partial_t \rho(x, t) dx \\ &= \int_x \mu(x)(1 - \rho(x, t)) \int_y \rho(y, t) w_\beta([\phi], y, x, t) dy dx \\ &\quad - \int_x \mu(x)\rho(x, t) \int_y (1 - \rho(y, t)) w_\beta([\phi], x, y, t) dy dx \\ &= - \iint dx dy \frac{Z(x, y)}{4D(x)D(y) \cosh[\frac{\beta}{2}(\phi(x) - \phi(y))]}, \end{aligned} \quad (\text{D.5})$$

where we have replaced w_β by the logit decision rate [see Eq. (7.12)] and we have symmetrized and simplified the second line to obtain the third line, and where we define $Z(x, y) \equiv [\mu(y) - \mu(x)](e^{\frac{\beta}{2}(\mu(y) - \mu(x))} - e^{\frac{\beta}{2}(\mu(x) - \mu(y))}) > 0, \forall x, y$. We thus conclude that $\frac{d\mathcal{F}}{dt} \leq 0$, i.e. that $\mathcal{F}$ is a Lyapunov function of the hydrodynamic evolution when the utility is linear.

D.2 Local mean-field description and LSA

In this section, we consider a modified dynamics where agents are allowed to relocate on neighboring site only. For simplicity, we also consider that the system is one dimensional. It is thus possible to perform a Taylor expansion of the different fields assuming that all fields are smooth in the mean-field limit. The jump probability between two neighboring sites becomes

$$f_\beta[u(x+a) - u(x)] = f_\beta \left(a\partial_x u + \frac{a^2}{2}\partial_x^2 u \right), \quad (\text{D.6})$$

where a is the lattice size, and u is the utility on position x . The evolution of the density (for non-overlapping agents) is thus given by

$$\begin{aligned} \partial_t \rho &= \rho(x+a)[1 - \rho(x)]f_\beta(-a\partial_x u - \frac{a^2}{2}\partial_x^2 u) \\ &\quad + \rho(x-a)[1 - \rho(x)]f_\beta(a\partial_x u - \frac{a^2}{2}\partial_x^2 u) \\ &\quad - \rho(x)[1 - \rho(x+a)]f_\beta(a\partial_x u + \frac{a^2}{2}\partial_x^2 u) \\ &\quad - \rho(x)[1 - \rho(x-a)]f_\beta(-a\partial_x u + \frac{a^2}{2}\partial_x^2 u). \end{aligned} \quad (\text{D.7})$$

After Taylor expansion up to $O(a^2)$ and time rescaling, it turns out that the evolution equation simplifies into

$$\partial_t \rho = f_\beta(0) \partial_x^2 \rho - 2f'_\beta(0) \partial_x [\rho(1 - \rho) \partial_x u]. \quad (\text{D.8})$$

Then, expanding around an homogeneous state, we write $\rho = \rho_0 + \rho_1(x, t)$, $\phi = \rho_0 + \phi_1(x, t)$, and we obtain to leading order in the perturbation:

$$\partial_t \rho_1 = f_\beta(0) \partial_x^2 \rho_1 - 2f'_\beta(0) \rho_0 (1 - \rho_0) u'(\rho_0) \partial_x^2 \phi_1. \quad (\text{D.9})$$

In Fourier space the evolution of the mode k is given by $\partial_t \hat{\rho}_1 = \Lambda(k) \hat{\rho}_1$, with

$$\Lambda(k) = -k^2 f_\beta(0) \left(1 - 2 \frac{f'_\beta(0)}{f_\beta(0)} \rho_0 (1 - \rho_0) u'(\rho_0) \hat{G}_\sigma(k) \right). \quad (\text{D.10})$$

From this, we deduce that the homogeneous system is unstable if there exists a mode k^* such that

$$1 < 2 \frac{f'_\beta(0)}{f_\beta(0)} \rho_0 (1 - \rho_0) u'(\rho_0) \hat{G}_\sigma(k^*). \quad (\text{D.11})$$

This criterion is exactly the same as the one found for the non-local move dynamics.

D.3 Linear stability for two coupled populations

We consider the evolution of a perturbation of the homogeneous state in Eq. (7.42) (and in its coupled analogue for the field ρ_B). Close to the homogeneous state $\rho_A(x) \equiv \bar{\rho}_A$, $\rho_B(x) \equiv \bar{\rho}_B$, with $\rho_0 = \bar{\rho}_A + \bar{\rho}_B$, we expand the fields $\rho_Z(x, t) = \bar{\rho}_Z + \rho_{Z,1}(x, t)$, with $Z = A$ or B , and the perturbation fields are denoted with index 1. One also has $\rho(x, t) = \rho_0 + \rho_1(x, t)$, and $\phi_Z(x, t) = \bar{\rho}_Z + \phi_{Z,1}(x, t)$. Keeping leading order terms in Eq. (7.42) yields

$$\begin{aligned} \partial_t \rho_{A,1} = \Omega \omega_A \Big(& -\bar{\rho}_A f_{\beta_A}(0) \rho_1(x, t) - (1 - \rho_0) f_{\beta_A}(0) \rho_{A,1}(x, t) \\ & + 2(1 - \rho_0) \bar{\rho}_A f'_{\beta_A}(0) [\phi_{A,1}(x, t) \partial_1 u_A + \phi_{B,1}(x, t) \partial_2 u_A] \Big), \end{aligned} \quad (\text{D.12})$$

where $\partial_1 u_A$ is a shorthand notation for $\frac{\partial u_A}{\partial \bar{\rho}_A} [\bar{\rho}_A, \bar{\rho}_B]$. Taking the logistic function $f_{\beta_A}(0) = \frac{1}{2}$, $f'_{\beta_A}(0) = \frac{\beta_A}{4}$, the linear evolution simplifies into

$$\begin{aligned} \partial_t \rho_{A,1}(x, t) = \frac{\Omega \omega_A}{2} \Big(& -\bar{\rho}_A \rho_1(x, t) - (1 - \rho_0) \rho_{A,1}(x, t) \\ & + (1 - \rho_0) \bar{\rho}_A \beta_A [\phi_{A,1}(x, t) \partial_1 u_A + \phi_{B,1}(x, t) \partial_2 u_A] \Big). \end{aligned} \quad (\text{D.13})$$

Appendix D. Schelling model

Similarly, we obtain for the evolution of B :

$$\partial_t \rho_{B,1}(x, t) = \frac{\Omega \omega_B}{2} \left(-\bar{\rho}_B \rho_1(x, t) - (1 - \rho_0) \rho_{B,1}(x, t) \right. \quad (\text{D.14})$$

$$\left. + (1 - \rho_0) \bar{\rho}_B \beta_B [\phi_{A,1}(x, t) \partial_1 u_B + \phi_{B,1}(x, t) \partial_2 u_B] \right). \quad (\text{D.15})$$

Denoting $\hat{\rho}_Z(k, t)$ the Fourier transform of $\rho_{Z,1}(x, t)$, the evolution equation can be cast in Fourier space into

$$\partial_t \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix} = L \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix}, \quad (\text{D.16})$$

with

$$L = \frac{\Omega}{2} \begin{bmatrix} \omega_A(\bar{\rho}_B - 1 + (1 - \rho_0) \bar{\rho}_A \beta_A \hat{G}_\sigma(k) \partial_1 u_A) & \omega_A(-\bar{\rho}_A + (1 - \rho_0) \bar{\rho}_A \beta_A \hat{G}_\sigma(k) \partial_2 u_A) \\ \omega_B(-\bar{\rho}_B + (1 - \rho_0) \bar{\rho}_B \beta_B \hat{G}_\sigma(k) \partial_1 u_B) & \omega_B(\bar{\rho}_A - 1 + (1 - \rho_0) \bar{\rho}_B \beta_B \hat{G}_\sigma(k) \partial_2 u_B) \end{bmatrix}. \quad (\text{D.17})$$

For simplicity, we will consider that agents are equally rational ($\beta_A = \beta_B = \beta$) and that their moving rates are also identical ($\omega_A = \omega_B = \omega$).

We are looking for conditions to observe dynamical patterns and/or static phase separation. Notably, the homogeneous state is linearly unstable if one eigenvalue of L has a positive real part. It is important to stress that the linear stability analysis is unable to predict the dynamic behavior when nonlinear terms become relevant. Whether the eigenvalues display an imaginary part or not *does not* bring any information on the final dynamics of the system. For the sake of completeness, we explicitate the criteria to have eigenvalues with positive real part and zero imaginary part, referred to as case (i), and eigenvalues with positive real part and nonzero imaginary part, referred to as case (ii). We lie in case (i) if

$$\begin{cases} \text{Tr } L > 0 \\ (\text{Tr } L)^2 - 4 \det L > 0, \end{cases} \quad \text{or} \quad \begin{cases} \text{Tr } L < 0 \\ \det L < 0. \end{cases} \quad (\text{D.18})$$

Case (ii) is obtained if

$$\begin{cases} \text{Tr } L > 0 \\ (\text{Tr } L)^2 - 4 \det L < 0. \end{cases} \quad (\text{D.19})$$

The criterion $\text{Tr } L > 0$ notably simplifies into

$$\bar{\rho}_A \partial_1 u_A + \bar{\rho}_B \partial_2 u_B > \frac{1}{\beta \hat{G}_\sigma(k)} \left(\frac{2 - \rho_0}{1 - \rho_0} \right). \quad (\text{D.20})$$

In the main text we have come up with utility functions that lead to eigenvalues with positive real parts and non zero imaginary parts, thus suggesting chasing instability. In some cases, oscillations were observed close to the homogeneous state but they eventually vanished at late times. Whether or not the chasing instability or oscillations are sustained cannot be predicted from the simple linear stability analysis but would require to perform a weakly non-linear analysis which is beyond the scope of this present paper.

D.4 LSA for two populations with local moves

We start from the local jump approximation of the mean-field equation for the coupled fields. We find that the dynamics can be cast into

$$\partial_t \rho_A = \partial_x [\rho_A (1 - \rho_A - \rho_B) \partial_x \mu([\rho_{A,B}], x)], \quad (\text{D.21})$$

with $\mu = \mu_{\text{ent.}} + \mu_{\text{util.}}$,

$$\mu_{\text{ent.}} = w_{\beta_A}(0) \log \left(\frac{\rho_A}{1 - \rho_A - \rho_B} \right), \quad (\text{D.22})$$

$$\mu_{\text{util.}} = -2w'_{\beta_A}(0) u_A([\rho], x), \quad (\text{D.23})$$

and likewise for ρ_B . One can look into the stability of an homogeneous state with densities $\bar{\rho}_A$ and $\bar{\rho}_B$, expanding around this state with a utility $u(\phi_A, \phi_B)$ for agents A and $v(\phi_A, \phi_B)$ for agents B . For convenience, we will take $w_{\beta_A}(0) = w_{\beta_B}(0) = \omega f_\beta(0) = \omega/2$ and $w'_{\beta_A}(0) = w'_{\beta_B}(0) = \omega f'_\beta(0) = \omega\beta/4$. Expanding around the homogeneous state $(\bar{\rho}_A, \bar{\rho}_B)$ leads to

$$\begin{cases} \partial_t \rho_{A,1} = \frac{\omega}{2} [(1 - \bar{\rho}_B) \partial_x^2 \rho_{A,1} + \bar{\rho}_A \partial_x^2 \rho_{B,1} - \partial_x (\beta \bar{\rho}_A (1 - \rho_0) \partial_x \phi_{A,1} \partial_1 u + \partial_x \phi_{B,1} \partial_2 u)] \\ \partial_t \rho_{B,1} = \frac{\omega}{2} [(1 - \bar{\rho}_A) \partial_x^2 \rho_{B,1} + \bar{\rho}_B \partial_x^2 \rho_{A,1} - \partial_x (\beta \bar{\rho}_B (1 - \rho_0) \partial_x \phi_{A,1} \partial_1 v + \partial_x \phi_{B,1} \partial_2 v)] \end{cases}, \quad (\text{D.24})$$

Hence, in Fourier space, the linear system can be cast into

$$\partial_t \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix} = K \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix}, \quad (\text{D.25})$$

with

$$K = \frac{\omega k^2}{2} \begin{bmatrix} \bar{\rho}_B - 1 + \beta \bar{\rho}_A (1 - \rho_0) \hat{G}_\sigma(k) \partial_1 u & -\bar{\rho}_A + \beta \bar{\rho}_A (1 - \rho_0) \hat{G}_\sigma(k) \partial_2 u \\ -\bar{\rho}_B + (1 - \rho_0) \bar{\rho}_B \beta \hat{G}_\sigma(k) \partial_1 v & \bar{\rho}_A - 1 + (1 - \rho_0) \bar{\rho}_B \beta \hat{G}_\sigma(k) \partial_2 v \end{bmatrix}. \quad (\text{D.26})$$

It is interesting to note that the evolution matrix K is directly proportional to L and, as a consequence, the stability criterion of the homogeneous state with local moves is exactly the same as the one found for non-local moves.

Titre : Naviguer dans la complexité radicale : l'influence du désordre, d'une dynamique non-relaxationnelle et de l'apprentissage sur la coordination agrégée

Mots clés : Physique statistique, systèmes désordonnés, hors équilibre, transitions de phase, auto-organisation, systèmes complexes

Résumé : Comment un comportement collectif peut-il émerger d'un système radicalement complexe ? En contradiction avec le paradigme de l'"homo economicus" de l'économie classique, dans lequel un état stationnaire unique et optimal est supposé atteint par des agents rationnels, la théorie des verres de spin a démontré que des problèmes d'optimisation simples peuvent produire un nombre immense de solutions lorsque qu'ils incluent des interactions hétérogènes. L'objectif de cette thèse est d'étudier comment des agents peuvent se coordonner dans un environnement radicalement complexe.

Nous commençons par illustrer la notion de complexité radicale dans le contexte socio-économique en étudiant l'abondance de solutions dans un problème d'optimisation de portefeuille. Nous observons un très grand nombre de solutions quasi-optimales, et par conséquent une extrême sensibilité de la solution optimale aux paramètres du problème, remettant en question l'idée même de décision rationnelle. Par ailleurs, ces résultats peuvent être directement appliqués aux équilibres écologiques des équations de Lotka-Volterra multispécifiques avec autorégulation. Ayant motivé la nécessité d'une rationalité limitée dans la modélisation des agents, nous revisitons un objet classique en microéconomie reliant la demande aux changements de prix : la matrice de Slutsky. Cet exemple nous permet d'illustrer à la fois l'efficacité de la mécanique statistique à l'équilibre pour aller au delà de certaines croyances conventionnelles, mais aussi ses limites dans la modélisation d'agents supposés égoïstes.

Nous proposons ensuite un "jeu" inspiré du verre de spin de

Sherrington et Kirkpatrick, mêlant rationalité bornée, interactions non réciproques et apprentissage. Nous montrons que ce modèle présente une multitude de phénomènes hors d'équilibre, tels que des cycles limites, du chaos et du vieillissement. Pour comprendre ces observations, qui démontrent que l'apprentissage est insuffisant pour surmonter la complexité radicale, nous étudions l'explosion du nombre de points de points fixes du problème, puis sa dynamique non-relaxationnelle à l'aide de la Théorie du Champ Moyen Dynamique. En raison de la richesse phénoménologique qu'il présente, nous estimons que ce "SK-game", incarne un modèle unificateur pour les systèmes socio-économiques désordonnés.

La dernière partie de la thèse est consacrée à des systèmes hors d'équilibre exempts de désordre. Nous revisitons d'abord un modèle d'occupation de type Schelling, dans lequel des agents en interaction peuplent un réseau. En remplaçant les quartiers prédéfinis par une perception locale de la densité, le système est conduit hors d'équilibre, mais nous montrons que sa phénoménologie demeure largement inchangée et est relativement indépendante de la règle de décision choisie. Enfin, nous examinons l'influence plus générale des courants hors d'équilibre sur la relaxation vers des états stationnaires préservant une distribution d'équilibre donnée. En établissant des bornes sur le temps de corrélation d'observables quelconques, nous étudions l'(in)efficacité de certains courants en vue d'accélérer la relaxation du système vers son état stationnaire, offrant des perspectives pour évaluer la robustesse de certains phénomènes collectifs à d'autres règles de décision.

Title : Navigating radical complexity: the influence of disorder, nonrelaxational dynamics and learning on aggregate coordination

Keywords : Statistical physics, disordered systems, out-of-equilibrium, phase transitions, self-organization, complex systems

Abstract : How does aggregate coordination occur in complex systems? At odds with the conventional "homo economicus" paradigm of classical economics, in which a unique and optimal steady-state is assumed to be reached by rational agents, spin-glasses have demonstrated that simple problems may yield an overwhelming number of solutions when heterogeneous interactions are present. The aim of this thesis is to study how idealized agents may navigate the radical complexity that likely emerges from their environment.

In a first part, we start by illustrating the notion of radical complexity in the socioeconomic context by studying the abundance of solutions in a constrained portfolio optimization problem. We find that there is a very large number of quasi-optimal solutions, and therefore an extreme sensitivity of the best possible outcome to the specific draw of problem parameters, challenging the very notion of a rational decision in this context. Remarkably, these results may be directly applied to ecological equilibria of multispecies Lotka-Volterra equations with self-regulation. Having motivated boundedly rational decision making in the modeling of individual agents, we revisit a classical object of consumer choice theory measuring the sensitivity of demand following a price change: the Slutsky matrix. This example allows us to illustrate both the effectiveness of equilibrium statistical mechanics to go beyond some conventional beliefs, but also its limitations to model agents that are assumed to be individualistic.

We then introduce a game inspired by the Sherrington-Kirkpatrick spin-glass, combining bounded rationality, non-reciprocal interactions and learning. We show that this model displays a wealth of dynamical phenomena, such as limit cycles, chaos and aging. To understand these observations, which demonstrate that learning is insufficient to overcome radical complexity, we study the explosion of the number of fixed points solutions to the problem, and its nonrelaxational dynamics with Dynamical Mean-Field Theory. Due to the very rich phenomenology it offers, we believe that this "SK-game" offers a unifying model for disordered socioeconomic systems.

A last part is devoted to disorder-free out-of-equilibrium systems. We first revisit a Schelling-like occupation model, in which interacting agents populate a lattice. While replacing fixed neighborhoods by a space-dependent perception of the density drives the system out of equilibrium, we find that its phenomenology is largely unchanged and independent of the details of the decision rule. We finally consider the more general influence of out-of-equilibrium forcing on the relaxation towards steady-states that have the particularity of matching a prescribed equilibrium distribution, allowing us to isolate the role of irreversible currents. By deriving upper bounds on the correlation time of arbitrary state variables, we study how given out-of-equilibrium contributions may (or may not) effectively accelerate the relaxation of the system, which could in turn be useful to assess the robustness of some collective phenomena to a wide range of decision rules.