

Paolo Galeazzi



Play Without Regret

Play Without Regret

Paolo Galeazzi

Play Without Regret

ILLC Dissertation Series DS-2017-01



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: <http://www.illc.uva.nl/>

The investigations were supported by the European Research Council under the European Research Community's Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement no.283963.

Copyright © 2017 by Paolo Galeazzi

Cover design by Paolo Galeazzi. Picture by Diego Fedele.
Printed and bound by Ipskamp Drukkers.

ISBN: 978-94-028-0463-8

Play Without Regret

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Aula der Universiteit
op woensdag 25 januari 2017, te 11.00 uur

door

Paolo Galeazzi

geboren te Venetië, Italië

Promotiecommissie

Promotor:	Prof. dr. S.J.L. Smets	Universiteit van Amsterdam
Co-promotor:	Dr. A. Baltag	Universiteit van Amsterdam
Overige leden:	Prof. dr. M. van Lambalgen	Universiteit van Amsterdam
	Prof. dr. R.A.M. van Rooij	Universiteit van Amsterdam
	Prof. dr. J. van Eijck	Universiteit van Amsterdam
	Dr. J. Szymanik	Universiteit van Amsterdam
	Prof. dr. R. Bradley	London School of Economics and Political Science
	Dr. K.J.S. Zollman	Carnegie Mellon University
	Dr. A. Perea	Universiteit van Maastricht

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

The results in this thesis come from the following works.

1. Chapter 4 is based on: Galeazzi, P. and Lorini, E. (2016). Epistemic logic meets epistemic game theory: a comparison between multi-agent Kripke models and type spaces. *Synthese*, 193 (7):2097-2127.
2. Chapter 5 and part of Chapter 7 are based on: Galeazzi, P. and Franke, M. (2016). Smart Representations: Rationality and Evolution in a Richer Environment. To appear in *Philosophy of Science*.
3. Chapter 6 is based on ongoing research with Mathias W. Madsen.
4. Chapter 8 is based on ongoing research with Zoi Terzopoulou.

*to everyone I have neglected
while I was busy with this dissertation (myself included)*

Contents

Acknowledgments	xv
1 Introduction	1
1.1 Subjective Expected Utility Maximization	3
1.2 Criticisms of SEU	4
1.3 A Brief Overview	6
1.4 Our Contributions	7
I Background	9
2 Background on Rational Choice	11
2.1 From the Early Days to Savage	11
2.2 Non-Bayesian and Non-Expected Utility Theories	19
2.3 Ecological Rationality	39
3 Background on (Evolutionary) Game Theory	43
3.1 Game Theory	43
3.2 Evolutionary Game Theory	58
3.3 An Important Example	68
4 Epistemic Structures for Rationality	75
4.1 Introduction	75
4.2 Models With Belief	77
4.2.1 Language	77
4.2.2 Semantics	77
4.2.2.1 Doxastic models	77
4.2.2.2 Type spaces	79
4.2.3 Correspondence	81

4.3	Models With Belief and Knowledge	84
4.3.1	Language	84
4.3.2	Semantics	84
	4.3.2.1 Type spaces	84
	4.3.2.2 Epistemic-doxastic models	85
4.3.3	Correspondence	86
4.3.4	Examples	88
4.3.5	Discussion	89
	4.3.5.1 On the knowledge operator	90
	4.3.5.2 Qualitative vs. quantitative type spaces	92
4.4	Axiomatization	93
4.5	Probabilistic Extension	99
4.5.1	Language	99
4.5.2	Semantics	100
	4.5.2.1 Type spaces	100
	4.5.2.2 Epistemic-probabilistic models	100
4.5.3	Correspondence	101
4.6	Conclusion	102

II Evolutionary Analysis 105

5	The Game of Life 107
5.1	The Behavioral Gambit 107
5.2	Choice Principles: A General Discussion 108
5.3	On the Rationality of Choice Principles 113
5.4	The Game of Life 114
5.5	The Model 115
5.6	A Bold Alternative to the Alternatives 116
5.7	Results 119
	5.7.1 Simulation results 119
	5.7.2 Analytical results 121
5.8	Extensions 122
	5.8.1 More preference types 122
	5.8.2 More actions 123
	5.8.3 Solitary decisions 125
5.9	Conclusion 126
	5.9.1 Subjective preferences or choice principles? 126
	5.9.2 Precise and imprecise beliefs 127
5.10	Appendix: Proofs 128

6	Learning Players' Types	137
6.1	The Learning Model	138
6.2	Intermezzo: A Simple Extraction	140
6.3	The Game of Life with Incomplete Information	141
6.4	Symmetry and Learning	144
6.5	Strong Informativity	145
6.6	The Set of Strongly Informative Games	149
6.7	Asymmetric Equilibria and Behavioral Assumptions	155
6.8	Discrete and Continuous Fitness	159
6.9	Conclusion	160
6.10	Appendix: Proofs	160
7	Evolution and Learning	167
7.1	The Game of Life with Learning	167
7.2	Play Without Regret	168
7.3	Regret What?	169
7.4	Conclusion	171
7.5	Appendix: Proofs	171
8	A Theory of Rationality	179
8.1	The Traveler's Dilemma Case	179
8.2	A Modest Proposal	180
8.3	Relief Maximization	184
8.4	From an Evolutionary Point of View	188
8.5	Nested and Weighted Principles	189
8.6	The Power of Context Dependency	191
8.7	Conclusion	195
9	Conclusion	197
	Bibliography	201
	Abstract	213
	Samenvatting	215

Acknowledgments

First of all, I want to thank my supervisors, Sonja Smets and Alexandru Baltag, for their valuable suggestions and their support along these years. The complete trust that I received from Sonja and Alex in developing my research directions is still a mystery to me, but it has certainly been a key factor for this dissertation. I hope I have been able to use their trust and their support at best.

Second, I am very grateful to Michiel van Lambalgen, Robert van Rooij, Jan van Eijck, Jakub Szymanik, Richard Bradley, Andrés Perea, and Kevin Zollman for accepting to be members of my thesis committee, which involves the heavy task of reading and understanding my exotic English sentences. In addition, I consider all the exchanges that I had with them extremely pleasant and encouraging from a personal point of view, and greatly profitable for my research.

The list of persons that I met during the Ph.D. and that contributed to the development of the ideas which gave birth to this thesis is still very long. I am surely the main responsible for the final outcome, and, hence, the first person to blame. However, the contribution of Michael Franke has been essential in these years. Coming from different fields, when I proposed him to work on my strange intuitions about the evolution of choice principles, he had absolutely no need or obligation (and probably even no much time) to give me any attention. I consider the fact that we immediately started working together not only the turning point of my Ph.D., but also a unique personal pleasure. It is likely that most of my actual research wouldn't even have begun without Michael.

Mathias Madsen has also importantly contributed to the improvement of this thesis. It was my great pleasure to develop the parts about the learning with him, we ended up learning a lot together, and I believe his brightness clearly shines through those pages.

I want to greatly thank Emiliano Lorini for inviting me at the University of Toulouse. That visit allowed me to enhance the research on type spaces and Kripke models which I had started for my master thesis, and which is contained here in Chapter 4.

I am also very thankful to Zoi Terzopoulou, who had the insane bravery to trust me and to ask for a project on decision theory. That project gave me the opportunity to develop some ideas on context-dependent rationality that I would have certainly forgotten otherwise. She has been a very smart and interested student, and I hope she had as much fun as I had during the project.

The outcome of my research would have also been considerably different without the many discussions that I had the privilege to entertain with all the persons that shared their time and their views with me along these four years. In this respect, I am especially indebted to Pierpaolo Battigalli, Erik Quaeghebeur, Teddy Seidenfeld, and Michiel van Lambalgen. Moreover, I am deeply grateful to Adam Brandenburger, Branden Fitelson, Joe Halpern, Massimo Marinacci, Eric Pacuit, and Peter Wakker for their attention and their valuable advices.

It is sometimes said that friends are the family that you choose for yourself. Coming from decision theory, I generally tend to ponder a lot and to choose carefully, but I think I have also been extremely lucky in human relations, and I am immensely thankful to all the persons that have been part of my life. Thanks to Michele, for our timeless friendship, and for his relentless effort for a better world, which makes me still believe that such a world is possible. Thanks to Hugo and Rasmus, I am honoured to have them as paronyms for my thesis defence, and, most importantly, as friends in life. The time we spent together will always be the dearest memory of these years. Thanks to Johannes and Riccardo, for being such wonderful flatmates and colleagues. Thanks to Thomas, Phaedon, Bastiaan, Dominik, Carlos, Nadine, Giovanni, Dieuwke, Malvin, Frederik, Ana Lucia, Ronald, Iris, Arnold, Julia, Ivano, Kostis, Pablo, Fernando, Nigel, Chenwei, Soroush, for the help, the warmth, the football, the music, the travels, the beers, the kindness, the lunches, the coffees, and in general for making me feel at home during the cold winters and the cold summers in this northern country, which in the end I loved especially because of them. Thanks to Matteo, Giorgio, Costanza, Valerio, Tommaso, Pedro, Andrea, Carlo, Jacopo, Samuele, Daria, Marce, Pau, and all the others, for always remembering about an old friend who emigrated years ago. Although we are often far apart, the moments we share in the same place make me feel like we have always been together. Special thanks again to Tommaso, Thomas, and Jean for improving this dissertation with their attentive observations and insightful comments. Thank you all for being such good friends, I hope I have been a good one too.

Finally, if in a life, other than friends, one should also choose his or her own family, there is no doubt that I would choose exactly the same, my mother Laura, my father Giorgio, and my brothers Alessandro and Stefano. They are the best family one could hope for. This thesis is dedicated to them.

Chapter 1

Introduction

What men really want is not knowledge but certainty. (B. Russell)

A basic distinction in scientific research is between *descriptive*, or positive, and *normative* sciences. The purpose of the descriptive approach is clear and well-defined: the definition of models that can describe actual observed phenomena. The meaningfulness of a normative approach, though, is less uncontroversial.

Generally speaking, we could say that the goal of normative sciences is to state how things *should* be, as opposed to positive sciences that aim at describing how things *are*. In some fields, such as physics or biology, it would be hardly sensible to hold any normative stance. The behavior of photons and electrons is just as it is, and there is no reasonable claim that it should be otherwise. However, normative questions fundamentally arise in social sciences, like ethics, economics and philosophy. For instance, defending a certain moral claim amounts to arguing about how people's behavior *should* be.

The present research is centered around the concept of rationality, and the issue of rational choice under uncertainty in particular. Although we all aim at choosing rationally throughout our everyday life, it seems that we make our choices without any established principle of rationality at hand. When we say that some action was (or was not) rational, we often appeal to intuitive and undefined insights about what it could be justified as the smart thing to do. But, of course, different actions may be justified on different bases, which complicates the problem of determining how one should act in a given decision situation. Just as a wayfarer that has no fixed star to follow and decides at any crossroad according to the spur of the moment, social sciences haven't yet identified the principle that serves us as the fixed guide in our everyday decisions.

Nowadays, the issue of rational choice is mainly addressed by economists, psychologists and philosophers, and the state of the art is far from reaching an agreement on what it means to choose rationally. Our investigation will take

place primarily in the fields of decision theory and game theory, as those parts of economics and mathematics that study individual and interactive decision making. Decision and game theory, however, do not necessarily have to deal with the issue of rationality. Both also admit a descriptive perspective, with the aim of modeling people’s actual behavior in (interactive) decision situations: this is the direction taken by behavioral economics.

In line with other prominent economists, rationality is intended here as a normative notion. We could then say that rationality defines the way an agent should choose, which opens to two possible readings of the term. An objective viewpoint would advocate that there are choices that are objectively better. Loosely speaking, one may be tempted to claim that, in general circumstances and apart from concocted exotic cases, not committing suicide is objectively better than committing suicide, and that choosing to commit suicide is objectively irrational. This approach may be related to and reminiscent of the concept of ecological rationality and the work done by Herbert Simon (e.g., [Simon, 1955], [Simon, 1990], [Simon, 1992]), and Gerd Gigerenzer and colleagues (e.g., [Gigerenzer and Brighton, 2009], [Gigerenzer, 2008]). From a subjective perspective, instead, a decision is rational if it is the best decision in the agent’s own eyes.¹ Rationality thus corresponds to what an agent should choose in the sense of how the agent would like to choose, or would like to have chosen. In this sense, an action is rational if the agent does not feel embarrassed about his or her decision and would not want to choose differently after a further analysis of the problem. This position is held in particular by [Gilboa, 2015]. In the following we will often return to both the more objective and ecological view of Simon and Gigerenzer, and the more subjectivist standpoint of Gilboa, Postlewaite and Schmeidler.² These introductory sections are supposed to start defining and clarifying our position on the matter, which is sometimes closer to the view of Gilboa, Postlewaite and Schmeidler, and other times is more in line with the approach taken by Gigerenzer and colleagues.

While the purpose of positive sciences is to describe reality, the goal of normative sciences is to *change* reality. Dealing with rationality on a normative ground means arguing for some specific decision making processes that a good decision maker should adopt, just as normative ethical claims should induce the agents to adhere to the underlying moral principles. An alternative account would be

¹A similar distinction can also be found in [Gilboa et al., 2010], where the authors axiomatize two different notions of rationality.

²When referring to the group of authors Gilboa, Postlewaite, and Schmeidler, we point to a series of different works, some of which are not authored by the three of them together. These works may present slightly different claims, but since they are unified by the same general spirit we allow ourselves to treat them as representative for a unified perspective on the topic. By the same reasoning, many papers falls under the label “Gigerenzer and colleagues”. Given the general ideas that they all share, we will tend to consider them together, as different developments of the same research direction.

to consider and simply describe the rationale behind the choices of the decision maker, just as a descriptive ethics would simply try to list the moral principles at the basis of the agent's behavior. To repeat, this research aims instead at investigating which choice principles decision makers should follow and which behavior they might be willing to exhibit.

1.1 Subjective Expected Utility Maximization

The view that has been dominant in economics for the last five decades relates the concept of rationality to the maximization of subjective expected utility. As we will see in detail, this paradigm traces back to the work by [Ramsey, 1926] and [de Finetti, 1931] on one side, and by [von Neumann and Morgenstern, 1944] on the other, and culminates in the axiomatization given by [Savage, 1954]. In short, this school of thought holds that a choice is rational if and only if it maximizes subjective expected utility, that is: the decision maker should have (i) a probabilistic belief about the possible states of the world, as well as (ii) a preference over the consequences that she can attain through her decision which can be numerically expressed by a subjective utility function over outcomes, and should finally make the choice that maximizes her utility with respect to her belief.

The main reason why subjective expected utility (SEU, hereafter) maximization became the paradigm of rational choice consists in the axiomatization offered in [Savage, 1954]. Savage's theorem is a jewel in itself and represents one of the best achievements ever reached in mathematical economics: he was able to present a set of seven axioms on a binary preference relation over the decision maker's possible options and to prove that an agent satisfies the axioms if and only if she can be represented as maximizing a subjective utility for a certain subjective probabilistic belief. Savage's axioms do not only look extremely elegant and logically compelling, they also offer a straightforward equivalence between the abstract notion of subjective expected utility maximization and the observable choices of a decision maker. In times when the positivist prescription of reducing scientific notions to measurable observables was the fundamental approach in philosophy of science, the latter played a crucial role for the success of Savage's theory. However, Savage's axioms have also important drawbacks, which immediately challenged the identification of rationality with SEU maximization.

SEU maximization corresponds to a concept of rationality that is, in some respects, rather weak. The standard for rational choice delivered by SEU reduces the essence of rationality to a matter of internal consistency. SEU never questions the particular preferences and beliefs of the decision maker: being consistent with your utility and belief (whatever they are) is all that is required to be rational, and all subjective utilities and probabilistic beliefs are considered equally rational by the theory. At the same time, though, SEU maximization may be seen as

expressing too strong a standard for rational choice. Both these directions of criticism are expanded and developed in the next section.

1.2 Criticisms of SEU

SEU maximization is often referred to as the Bayesian paradigm in decision theory. Specifically, according to [Gilboa et al., 2012], the Bayesian paradigm in decision theory consists of four main tenets:

1. The Grand State Space;
2. Prior probability;
3. Bayesian updating;
4. Expected utility maximization.

It is a common view in microeconomics that a state of the world should resolve all uncertainty. Consequently, all the parameters that are relevant to the economic agent must be specified in each possible state. With time, from Savage to Harsanyi ([Harsanyi, 1967], [Harsanyi, 1968a], [Harsanyi, 1968b]) and to Aumann ([Aumann, 1976]), this view has enlarged the amount of information needed to be specified in a state, and caused an exponential growth in the number of possible states. The final outcome of this process is what has been called the *Grand State Space* (e.g., [Gilboa et al., 2012]). According to Gilboa, Postlewaite and Schmeidler, the assumption of the Grand State Space may already be problematic, especially in combination with the second tenet: in some cases, holding a probability measure over such a space can be hardly feasible for the decision maker.³

The second tenet states that all the agent's uncertainty must always be quantified by means of a probability measure over the state space. This requirement has firstly come under attack from a descriptive standpoint. Ellsberg's paradoxes (see Chapter 2) became paradigmatic examples of such a criticism.

The third tenet, namely that agents should update their prior probability according to Bayes' rule, has been considered the most uncontroversial assumption from a normative point of view, although the claim that agents actually follow Bayesian updating has also been disproven by empirical data (see, for instance, [Tversky and Kahneman, 1974]).

Finally, it is assumed in the Bayesian paradigm that rational economic agents choose the option that maximizes expected utility, given their subjective preferences and subjective beliefs. The fourth tenet has also been severely confuted by

³For more criticism of the Grand State Space assumption, and the combination of the first two tenets of the Bayesian paradigm, see [Gilboa et al., 2009], [Gilboa et al., 2012], [Gilboa and Schmeidler, 1995], and [Gilboa and Schmeidler, 2001].

many psychologists, economists and philosophers. The renowned Allais' paradox (see Chapter 2) made immediately evident that decision makers' choices aren't always compatible with expected utility maximization. The work by Daniel Kahneman and Amos Tversky subsequently documented a plethora of experimental situations in which observed behavior violates not only the tenet of expected utility maximization, but almost all of the axioms and assumptions of SEU theory (e.g., [Kahneman and Tversky, 1979]). Famously, Amos Tversky used to say "Give me an axiom and I'll design the experiment that refutes it." ([Gilboa, 2010]).

Although SEU theory has been confuted from a descriptive point of view by empirical evidence, it is the normative appeal of its axioms that still preserves it as the golden standard of rational choice. The arguments advanced by Gilboa, Postlewaite and Schmeidler, instead, aim at rejecting subjective expected utility maximization on a normative ground. Their work mainly focuses on the issue of *rationality of beliefs*, and leaves aside questions about the rationality of the two other components of the theory, that is, subjective preferences and expected utility maximization as the decision criterion.⁴ They justify their approach by the fact that "one may say more about beliefs than about tastes." Indeed, as they continue in [Gilboa et al., 2012]: "Rationality does not constrain one to like or to dislike the smell of tobacco but rationality does preclude the belief that smoking has no negative health effects." They support a higher standard of rationality, that implies the rationality of beliefs: beliefs should be justified by evidence. In their opinion, the Bayesian paradigm is lacking a theory of belief formation and a classification of beliefs according to their rationality:

A theory of belief formation could suggest a systematic way of predicting which beliefs agents might hold in various environments. [...] In particular, we would be able to tell when economic agents are likely to entertain probabilistic beliefs, and when their beliefs should be modeled in other, perhaps less structured ways.

This shortcoming makes SEU too weak and too strong at the same time. It is too weak for the already mentioned reason that whatever belief, if probabilistic, is considered legitimate and rational. It is too strong because it does not allow any uncertainty representation different than a probabilistic belief. In line with the higher standard of rationality that they want to sustain, [Gilboa et al., 2012] argue that: "Justification of beliefs by evidence offers a criterion for rationality that need not rank highly specified beliefs as more rational than less specified ones."

We fully agree with Gilboa, Postlewaite and Schmeidler on this point, and, to conclude the *pars destruens*, we finish with a last quote from [Gilboa, 2015] which lucidly sums up one of the major difficulties that we also see in the Bayesian paradigm:

⁴Similar positions can also be found, among others, in [Levi, 1974] and [Gärdenfors and Sahlin, 1982].

The Bayesian approach is quite successful at representing knowledge, but rather poor when it comes to representing ignorance. When one attempts to say, within the Bayesian language, ‘I do not know’, the model asks, ‘How much do you not know? Do you not know to degree .6 or to degree .7?’ One simply doesn’t have an utterance that means ‘I don’t have the foggiest idea’.

It is therefore not at all obvious that rationality – even subjective rationality – suggests that we select one prior out of all possible ones. [...] Many would probably agree that they would feel more comfortable with a choice of a paradigm that can represent ignorance as well as knowledge.

The *pars construens* must then consist of possible alternatives for decision making where the representation of uncertainty need not be probabilistic. Two proposals along these lines come in particular from the work done by Gilboa and Schmeidler in the 1980s. Specifically, [Schmeidler, 1989] gives axiomatic representation of a decision making based on non-additive probability measures and Choquet expected utility. [Gilboa and Schmeidler, 1989] instead axiomatize *maxmin expected utility* decision criterion with non-unique prior, i.e., situations where the decision maker acts as if she held a (convex compact) set of prior probabilities and picked the option that maximizes the minimal expected utility, ranging over all the priors in the set. From a descriptive point of view, both these alternatives would for example not rule out behaviors such as those observed in Ellsberg’s paradoxes, but they primarily stemmed from normative, rather than descriptive, considerations about the excessive limitations imposed by the Bayesian school. A detailed exposition of the relevant decision-theoretic literature, including Savage’s and Gilboa and Schmeidler’s axiomatizations as well as the presentation of Allais’ and Ellsberg’s paradoxes, is introduced in Chapter 2.

1.3 A Brief Overview

In the chapters of Part II, we build on the (normative) assumption that not all uncertainty can be quantified in a probabilistic fashion, and we try to challenge the rationality of the two other main components of SEU: subjective preferences and expected utility maximization as the decision criterion.

In particular, we consider more general decision criteria for situations where beliefs might be non-probabilistic (that would eventually reduce to expected utility maximization in case of probabilistic uncertainty), and we investigate the (ecological) rationality of different decision criteria. For our purposes, the direction taken by [Gilboa and Schmeidler, 1989] will be especially relevant: when beliefs are represented in terms of a set of prior probabilities, maxmin expected utility is just one among many possible criteria. The approach of the present

research to the issue of the rationality of decision criteria is *ecological*, and it is close in spirit to the work of Herbert Simon and to the school of Gigerenzer and colleagues. The rationality of different decision criteria is studied and assessed in relation to the environment where decisions take place, that is also composed of *interactive* decision situations in which different criteria have to compete with each other. Game theory and, more specifically, evolutionary game theory will thus be the principal tool for this investigation. The necessary background in (evolutionary) game theory is given in Chapter 3.

On the other hand, subjective preferences suffer from the same weakness that we have seen for subjective beliefs: all subjective utility functions are equally rational. All that matters is that choices are consistent with a certain utility and a certain belief. In our opinion, questions about the rationality of preferences should also naturally arise here. For instance, when one is thirsty, it does not seem rational to have a preference for drinking mercury over drinking water. But on what basis would one consider it irrational? A purely subjectivist approach would try to defend the indisputability of individual tastes, but there is also a sense in which some preferences are better than others. This sense, we claim, is grounded in the notion of evolutionary fitness.

In fact, there is a research direction, recently developed in economics under the name of *evolution of preferences*, that aims at studying the evolutionary competition between different subjective utilities (e.g., [Samuelson, 2001], [Robson and Samuelson, 2011], [Dekel et al., 2007]). Evolution of preferences has already proven to be an insightful approach where the existence of apparently (from an individualistic point of view) irrational preferences can be explained on the basis of evolutionary and environmental factors (e.g., [Alger and Weibull, 2013]). Chapter 5, Chapter 6 and Chapter 7 present the main results about the ecological rationality of different decision criteria and subjective preferences.

To sum up, this study wants to challenge the dominant position that views subjective expected utility maximization as the norm for rational choice. To that aim, SEU theory is considered as the combination of three individual components: subjective preferences, subjective probabilistic beliefs, and expected utility maximization as the decision criterion. While normative arguments against the necessity of expressing uncertainty by means of a probability measure have already been introduced, we are going to question the ecological rationality of the other two constituents: subjective preference and expected utility maximization criterion. Finally, some guiding lines for a more realistic theory of rational choice are proposed in Chapter 8.

1.4 Our Contributions

This section is meant to highlight the main original contributions of this work, and to help the reader to go through them.

As already anticipated in the previous section, the main results about ecological rationality are developed in Chapter 5, Chapter 6, and Chapter 7. We refer the reader directly to those chapters for what we consider the pivot of this thesis, and we are not going to anticipate here the outcomes of the evolutionary competitions investigated there. However, we want to stress that we view the model used to study such an evolutionary competition as a fundamental contribution per se. Although similar approaches can be found in a few recent papers (e.g., [Zollman, 2008], [Bednar and Page, 2007], [Robalino and Robson, 2016]), the multi-game model of this thesis, that we call *the game of life*, had not been formulated in the full generality that reaches here, for the best of our knowledge. The novelty and the contribution of the approach in itself is that it allows to encompass a (possibly infinite) variety of different (interactive) decision situations, and to study the evolutionary competition of general ways of making choices in such a rich and extended environment. In our opinion, this approach to evolutionary game theory would merit closer attention in the future. As we are going to argue in more detail at the beginning of Chapter 5, the behavioral gambit embedded in single-game models might often be a limitation if we think of evolution as driven by a series of different interactions, rather than a fixed single one.

In turn, the construction of the model is carefully built from standard notions of decision theory and (evolutionary) game theory, in order to be presented as an incremental extension of well-established literature. Chapter 2 and Chapter 3 represent a thorough selection and original exposition of all the background concepts that are needed for the development of the game of life. To give an example that will be clearer from the reading of the chapters, the foundation of the game of life with learning of Chapter 6 comes from the notions of Bayesian game, ex ante and interim equilibrium presented in Chapter 3, and in particular from the case of incomplete information games under ambiguity introduced in Section 3.3. The reader that is not familiar enough with these notions and that will try to immediately jump to the results of Part II will incur the risk of missing substantial ingredients of the arguments. We therefore view the background chapters as important contributions to the story line of this work. This is obvious for Chapter 4 (that offers a logical analysis of the epistemic structures for reasoning about rationality used in epistemic game theory), but we also consider the presentation of Chapter 2 and Chapter 3 fundamental for the best understanding of all that is developed later.

Part I

Background

Chapter 2

Background on Rational Choice

Economists tend to think they are much, much smarter than historians, than everybody. And this is a bit too much because at the end of the day, we don't know very much in economics. (Thomas Piketty)

Formal decision theory started during the 16th and 17th century along with the pioneeristic works in probability theory by Pierre de Fermat, Blaise Pascal, and Christiaan Huygens. To this day, the literature on decision theory has been increasing exponentially, and is now so vast that an entire book would not be enough to exhaustively review it. This chapter does not aim to be a comprehensive overview of the research on rational choice, but rather to lay out the necessary notions and background for the developments of the following chapters.

2.1 From the Early Days to Savage

Blaise Pascal (1670) It is most likely with [Pascal, 1670]'s wager that many important ideas of modern decision theory made their first appearance. At that time, God was a major concern for scientists and philosophers. But while Decartes and Leibnitz were still seeking after a proof for God's existence, Pascal decided to address the issue from a different angle, and gave rise to the first modern formulation of a decision problem. Instead of the existence of God, Pascal's point was whether we should conduct a pious and religious life, or a worldly one. His framing of the problem is verbal, but it is clearly phrased in terms of a bet between two acts, as pictured in the table below.

	God exists	God does not exist
pious	eternal happiness	constrained life
not pious	sorrow	unconstrained life

Following Pascal’s argument, the verbal outcomes of the previous matrix may be expressed in util units as in the next table.

	God exists	God does not exist
pious	∞	b
not pious	c	d

The quantities b, c, d are such that: $d > b$ and $c > -\infty$. The bet proposed by Pascal is then dependent on the degree of belief that the agent has in the existence of God. Specifically, if we denote by p the agent’s probability assessment for the existence of God, the expectation from a pious life should be $p \cdot \infty + (1 - p) \cdot b$, whereas the expectation from a worldly life is $p \cdot c + (1 - p) \cdot d$. Pascal’s conclusion is therefore that, for any $p \in (0, 1]$, the agent should choose a pious conduct. Informally stated, if the agent does not fully exclude the possibility of God existing, and instead gives the existence some nonzero probability, then her best choice would be to behave in accordance with religious principles.

Remarkably, the behavioral shift from God’s existence to life conduct brought Pascal to introduce fundamental notions of theory of choice, such as the formalization of a decision problem by means of a decision matrix, the use of probabilities to express subjective degrees of belief, the representation of subjective utilities as numerical quantities, and the appeal to expected utility maximization for solving the decision problem.

Christiaan Huygens (1657) Apart from formulating the wave theory of light, discovering the first of Saturn’s moons, Titan, through a telescope he designed himself, inventing the pendulum clock and the pocket watch, and other things, the name of Christiaan Huygens also appears in the foundations of probability and decision theory. For our purposes, his main contribution is the explicit proposal of expected value as the proper criterion to evaluate games of chance, where probabilities are objectively given. In [Huygens, 1657], Huygens does not a priori assume that expected value is the suitable criterion to evaluate games of chance, but rather tries to demonstrate its adequacy in different games by mathematical proofs from more primitive principles. As pointed out by [Gilboa and Marinacci, 2013], “Huygens’ propositions can be thus viewed as the very first decision-theoretic representation theorems, in which the relevance of a decision criterion is not viewed as self-evident, but needs to be justified through logical arguments based on first principles.”

Daniel Bernoulli (1738) A step forward with respect to Huygens’ expected value criterion is marked by the work of Daniel Bernoulli, and in particular by his solution to the *Saint Petersburg Paradox* in [Bernoulli, 1738], a problem raised by his cousin Nicolas Bernoulli. The paradox goes as follows: the decision maker (DM hereafter) is offered a game of chance where a fair coin is tossed until it comes up Heads, and the prize for DM amounts to 2^n ducats, where the n^{th} toss

is the first time that Heads comes up. Before Daniel Bernoulli's contribution, the prevalent opinion about how lotteries and games of fortune had to be evaluated made appeal to the expected value criterion, as suggested by Huygens. When we compute the expected value of the Saint Petersburg Paradox, we get:

$$\sum_{n=1}^{\infty} 2^n \frac{1}{2^n} = 1 + 1 + 1 + \dots = \infty.$$

Consequently, according to the expected value criterion, a DM would be willing to pay an infinite amount (so, any amount) of ducats to enter the game. This conclusion is blatantly counterintuitive. Rather than looking at the objective value, the solution proposed by Daniel Bernoulli lies on the notion of subjective *utility*. As [Bernoulli, 1738] wrote,

But anyone who considers the problem with perspicacity and interest will ascertain that the concept of *value* which we have used in this rule may be defined in a way which renders the entire procedure universally acceptable without reservation. To do this the determination of the *value* of an item must not be based on its *price*, but rather on the *utility* it yields. The price of the item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate.

Bernoulli observed that the increase in subjective utility that DM experiences when an additional ducat is added to her amount must be supposed to decrease as her amount increases. This is the principle of *diminishing marginal utility*, which has been a central tenet of economic theory ever since. Formally, this intuition led Bernoulli to conjecture that utility must be a concave function of the amount of money m , and he proposed to express it by the logarithm $\log(m)$. With this interpretation, the expected utility of the coin game would be:

$$\sum_{n=1}^{\infty} \log(2^n) \frac{1}{2^n} = \log(2) \sum_{n=1}^{\infty} n \frac{1}{2^n} = 2 \log(2) < \infty.$$

Given m the wealth of DM, the fair price p of the game is computed by equating DM's current utility $\log(m)$ with the expected utility of her final wealth if she pays the amount p and enters the game, that is,

$$\log(m) = \sum_{n=1}^{\infty} \frac{1}{2^n} \log(m - p + 2^n).$$

Intuitively, DM will be willing to enter the game for any price that is lower than p , and will not enter the game for any price that exceeds p . For example, if we take DM with initial endowment of 200 ducats and natural logarithmic utility

function, then DM would accept to enter the game only for prices lower than 9 ducats.¹

Daniel Bernoulli's work represented a major breakthrough for the development of economic theory. He formally introduced the notion of subjective utility, and from then on *subjective expected utility* has replaced expected value as the standard criterion for evaluating games of chance.

Frank Plumpton Ramsey (1926) The modern representation of uncertainty in decision theory and game theory is grounded in the so-called Bayesian approach, and is heavily indebted to the work done in the 1920s and 1930s by Frank Plumpton Ramsey and Bruno de Finetti. [Ramsey, 1926] was the first to explicitly state two main principles that are at the basis of the Bayesian approach.

1. The only sensible measure of subjective degrees of belief in some event E is DM's *willingness to bet* on event E .
2. The only *consistent* betting behavior of DM is in accordance with the laws of probability theory.

The first principle endorses a behavioral definition of belief. Simple introspection is not reliable to discover DM's beliefs: the only way to measure DM's beliefs about a certain event is to check how heavy DM is willing to bet on that event. As Ramsey noted too, this achievement had already been established in worldly wisdom since long ago. The old technique of resolving a matter of contention by challenging the opponent part to bet on her beliefs has probably been used since the invention of bets themselves. More importantly, the appeal to betting behavior enables theories of belief to acquire an operational basis that allows to relate theoretical notions to empirical observations.

The second principle connects the consistency of such a betting behavior with the use of probability theory to form beliefs. By *consistent* behavior Ramsey means a behavior that cannot be exploited by Dutch books. Standard Bayesianism holds that any uncertainty about events or propositions must be quantified in terms of a probability distribution, and Ramsey's observation that betting behaviors at odds with probabilistic beliefs can be Dutch-booked was a central contribution to the rise of the Bayesian paradigm. The two principles constitute the core of de Finetti and Ramsey's subjectivist view of probability theory.

Bruno de Finetti (1931) It is likely that nobody tried to promote and to justify Bayesianism as *the* method to reason about uncertainty as did Bruno de Finetti. He gave many considerable contributions to the fields of probability and statistics, but his paper from 1931 is especially relevant for the present work. [de Finetti, 1931] first introduces a binary relation \succsim over an algebra Σ of events

¹This example is due to Laplace ([Laplace, 1814]).

defined on a state space S . The relation is interpreted as *qualitative probability*, with $E \succsim E'$ denoting that the event E is at least as probable than E' . The goal of [de Finetti, 1931] is to state the properties of \succsim to guarantee that for any events $E, E' \in \Sigma$ we can always find a (finitely additive) probability function P over Σ , such that

$$E \succsim E' \text{ if and only if } P(E) \geq P(E'). \quad (2.1)$$

The axioms for \succsim used by de Finetti are the following:

- d1. \succsim is a total preorder, i.e., complete and transitive.
- d2. For all $E \in \Sigma$, $S \succsim E \succsim \emptyset$.
- d3. Given any events $E, E', E'' \in \Sigma$ such that $E \cap E'' = \emptyset = E' \cap E''$, $E \succsim E' \Rightarrow E \cup E'' \succsim E' \cup E''$.
- d4. For each $n \geq 1$, there is a partition of S into n equally probable events $\{E_i\}_{i=1}^n$.

De Finetti was able to prove that if \succsim satisfies d1-d4, then there exists a probability function P such that

$$\text{if } E \succsim E' \text{ then } P(E) \geq P(E'). \quad (2.2)$$

But his interest was originally the question whether axioms d1-d3 suffice to prove the result in (2.1). Years later, [Kraft et al., 1959] demonstrated that they do not, and that d1-d3 are not even sufficient to prove (2.2). The most satisfactory solution to this problem was given by [Scott, 1964]², who was able to show that for \succsim to be realizable by a probability measure, it is necessary and sufficient that the conditions

- 1. \succsim is total,
- 2. $\forall E \in \Sigma, S \succsim E \succsim \emptyset$,
- 3. $\mathbb{1}_E + \mathbb{1}_{E_1} + \dots + \mathbb{1}_{E_n} = \mathbb{1}_{E'} + \mathbb{1}_{E'_1} + \dots + \mathbb{1}_{E'_n} \Rightarrow E \succsim E'$,

hold for all $E, E_1, \dots, E_n, E', E'_1, \dots, E'_n \in \Sigma$ whenever $E_i \succsim E'_i$ (with $1 \leq i \leq n$), and where $\mathbb{1}_E \in \{0, 1\}^S$ is the characteristic function of event $E \in \Sigma$.

Apart from later improvements of the results, the accomplishments of de Finetti mark one of the first attempts to associate axioms for a qualitative relation with a quantitative and numerical representation of that relation. Through this procedure he offered Savage the first qualitative axiomatization of subjective probabilities in terms of a binary relation over events.

²Although Dana Scott was himself not satisfied by the result in that one of the conditions is an algebraic sum of characteristic functions, and not a strictly Boolean condition.

John von Neumann and Oskar Morgenstern (1944)

Behavioral foundations are mathematical results, stating a logical equivalence between a list of preference conditions and a decision model. They show that the subjective parameters in the decision model are the relevant parameters for determining decisions. [Wakker, 2010]

The next landmark in the development of modern decision theory is due to the joint work of two scholars emigrated from Europe and based at Princeton University, John von Neumann and Oskar Morgenstern. The representation result they succeeded in proving involves a binary preference relation \succsim over lotteries, and provides the behavioral foundation of expected utility theory. Given a set X of possible outcomes, a *simple lottery* is a probability function with finite support over the possible outcomes. For $x, y \in X$, we denote by $L = \{x, p; y, (1-p)\}$ the lottery that gives outcome x with probability p and outcome y with probability $1-p$. It is worth stressing that the set X can be any set of “things”, not necessarily numerical quantities: x might be a cow and y a trip to Hawaii, for example. In [von Neumann and Morgenstern, 1944], the preference relation is assumed to satisfy the following properties over the space of all (simple) lotteries $\Delta(X)$ on X :

vNM1. *Total preorder*: \succsim is complete and transitive.

vNM2. *Independence*: For $L, L', L'' \in \Delta(X)$ and $\alpha \in (0, 1)$,

$$\text{if } L \succ L' \text{ then } \alpha L + (1-\alpha)L'' \succ \alpha L' + (1-\alpha)L''.$$

vNM3. *Archimedean*: For $L, L', L'' \in \Delta(X)$ such that $L \succ L' \succ L''$, it is always possible to find $\alpha, \beta \in (0, 1)$ such that

$$\alpha L + (1-\alpha)L'' \succ L' \succ \beta L + (1-\beta)L''.$$

The representation theorem of [von Neumann and Morgenstern, 1944] states that the following are equivalent.

1. \succsim satisfies axioms vNM1-vNM3.
2. There exists a function $u : X \rightarrow \mathbb{R}$ such that

$$L \succsim L' \text{ iff } \sum_{x \in \text{supp}L} u(x)L(x) \geq \sum_{x \in \text{supp}L'} u(x)L'(x),$$

where $L(x)$ is the probability given by lottery L to outcome x , and $\text{supp}L$ denotes the (finite) support of lottery L , that is, the finite set of outcomes that L gives nonzero probability. Moreover, the function u , often called *vNM utility function*, is unique up to positive affine transformations. Through the vNM utility u over outcomes we can then define the subjective utility of lotteries $\bar{u} : \Delta(X) \rightarrow \mathbb{R}$ such that $\bar{u}(L) := \sum_{x \in \text{supp}L} u(x)L(x)$.

We have seen that subjective expected utility had already appeared as the criterion to evaluate chance games and lotteries since the solution of the Saint Petersburg Paradox by Daniel Bernoulli, and had been very popular thereafter. But despite of its convenience, it was not clear why one would exclude other properties of the utilities' distribution, such as the variance for instance. The work by von Neumann and Morgenstern has the merit of founding expected utility maximization on solid behavioral basis, formulated in terms of axioms on DM's preference relation. Their representation theorem is a milestone for the theory of choice under *risk*, i.e., for those situations where the probabilities of outcomes are evidently and objectively given, and DM is not asked any subjective probabilistic assessment.

Leonard Jimmie Savage (1954) The culmination of Savage's research undoubtedly lies in the book *The Foundations of Statistics* ([Savage, 1954]).³ There, from an axiomatic system on the preference relation \succsim , Savage is able to derive subjective expected utility maximization for cases where neither utilities nor objective probabilities are given. This is a step further with respect to [von Neumann and Morgenstern, 1944], because the probabilities of the outcomes are now unknown, and DM does no longer face a decision problem under risk, but rather a choice under *uncertainty*. The representation theorem in [Savage, 1954] rests upon the combination of von Neumann and Morgenstern's derivation of expected utility and the axiomatization of subjective probabilities in [de Finetti, 1931]. Savage's work has been so influential that mainstream economics has been indissolubly tied up with the Bayesian paradigm and subjective expected utility maximization from then on, and his model is now the standard model of a Bayesian decision problem.

The beauty and success of Savage's framework lies in its simplicity. There are only two primitives: states of the world S (with characteristic element s), and outcomes X (with characteristic element x). An event $E \subseteq S$ is any subset of S , which is considered to be endowed with the algebra 2^S of measurable events. The object of DM's choice are acts. An *act* is defined as a function $f : S \rightarrow X$, and let \mathcal{F} denote the set of all possible acts: $\mathcal{F} := X^S = \{f | f : S \rightarrow X\}$. It is to be noticed that all the *constant acts* (i.e., acts that yield the same outcome in any state s) are in \mathcal{F} . Let then $x \in \mathcal{F}$ denote the constant act that yields outcome $x \in X$ in any possible state, that is, the act f such that $\forall s \in S f(s) = x$. Moreover, for event $E \subset S$ and acts $f, g \in \mathcal{F}$, we denote by f_E^g the act such that

$$f_E^g(s) = \begin{cases} g(s) & \text{if } s \in E \\ f(s) & \text{if } s \notin E \end{cases}$$

³The title reflects the influence of [Wald, 1950]. At that time, the main concern was as decision theoretic as it was statistical: the search for "the set of principles that [...] an objective statistical decision rule should satisfy", in Chernoff's words ([Bather, 1996]).

Finally, let $f \succsim_E g$ denote the fact that DM prefers f over g given event $E \subseteq S$. Formally, \succsim_E can be defined by modifying the two acts f and g in such a way that they agree on the complement $E^c := \{s \in S \mid s \notin E\}$, i.e.,

$$f \succsim_E g \text{ if } f_{E^c}^h \succsim g_{E^c}^h, \text{ for } h \in \mathcal{F}.$$

An event $E \subset S$ is then called *null* if, for every $f, g \in \mathcal{F}$, it holds that $f \sim_E g$. Within this set-up, the preference relation $\succsim \subseteq \mathcal{F} \times \mathcal{F}$ is supposed to satisfy the following axioms.

- P1. \succsim is a total preorder, i.e., complete and transitive.
- P2. For $E \subset S$, and $f, g, h, h' \in \mathcal{F}$, $f_{E^c}^h \succ g_{E^c}^h$ iff $f_{E^c}^{h'} \succ g_{E^c}^{h'}$.
- P3. For $f \in \mathcal{F}$, $x, y \in X$, and non-null event $E \subset S$, $x \succ y$ iff $f_E^x \succ f_E^y$.
- P4. For $E, E' \subset S$, $x, y, w, z \in X$ such that $x \succ y$ and $w \succ z$, $y_E^x \succ y_{E'}^x$ iff $z_E^w \succ z_{E'}^w$.
- P5. $\exists f, g \in \mathcal{F}$ such that $f \succ g$.
- P6. For $f, g, h \in \mathcal{F}$ with $f \succ g$, there exists a partition $\{E_1, \dots, E_n\}$ of S such that, for every $i \leq n$, $f_{E_i}^h \succ g$ and $f \succ g_{E_i}^h$.
- P7. For $f, g \in \mathcal{F}$, $E \subset S$, if $\forall s \in E$ $f \succsim_E g(s)$ then $f \succsim_E g$, and if $\forall s \in E$ $f(s) \succsim_E g$ then $f \succsim_E g$.

Savage proved that \succsim satisfies axioms P1-P7 if and only if there exist a non-atomic⁴ finitely additive probability measure P on $(S, 2^S)$ and a non-constant bounded function $u : X \rightarrow \mathbb{R}$ such that, for any $f, g \in \mathcal{F}$:

$$f \succ g \text{ iff } \int_S u(f(s)) dP(s) \geq \int_S u(g(s)) dP(s).$$

Moreover, P is unique, and u is unique up to positive affine transformations, as in [von Neumann and Morgenstern, 1944].

As [Gilboa and Marinacci, 2013] write, to this day Savage's theorem "is universally viewed as the most compelling reason to assume that rational choice necessitates Bayesian quantification of all uncertainty, that is, the reduction of uncertainty to risk." Its normative force is primarily due to the clear and intuitive appeal of its axioms. With a view to later developments, two last points deserve to be highlighted here. First, the rationality principle delivered by Savage's work is the *maximization of subjective expected utility*: DM should choose the act that maximizes her expected utility in light of her subjective beliefs. Maximization of

⁴A probability measure is non-atomic if, for each event E s.t. $P(E) > 0$, there exists an event $E' \subset E$ s.t. $P(E') > 0$.

subjective expected utility embodies a weak, internal notion of rationality, whose only requirement is DM's consistency of choices, without any demand for her choices to be motivated by reasonable utilities and sensible beliefs. Second, there is a historical anecdote that is very interesting for what we are going to see in later chapters. According to [Bather, 1996], the ideas that originated Savage's investigation came from a conversation that he had with Herman Chernoff. It is worth reporting Chernoff's words directly from [Bather, 1996]:

At one time Savage felt that he had resolved the choice of criterion problem in decision theory. In the decision theory approach there remained a question of how you select among the various decision rules when there is not a uniformly best choice (which is usually the case). Wald had tentatively suggested the minimax principle.

[...] Savage proposed that minimax regret was the resolution to the problem. When I played with that notion, I found that it failed to satisfy one of Arrow's requirements for a universal choice function. That was the principle of irrelevant alternatives. If you had the choice of a, b, c or d , you might decide that a is the best. However, if someone tells you d is not available, it may then turn out that among a, b and c , you prefer b . Minimax regret sometimes behaved this way, and that was a violation of this principle of irrelevant alternatives. I brought this to Savage's attention and, after arguing futilely for a little while that it proved how good his criterion was, he finally agreed that it was wrong. He felt then that perhaps we should be elaborating on de Finetti's Bayesian approach, which he had come across. (He was a voracious reader.) Meanwhile, I decided that I would list the set of principles that I felt an objective statistical decision rule should satisfy. I wrote a discussion paper on rational selection of decision functions which came up with a contradiction. I sat on it for a few years until I finally published it in *Econometrica*. Ultimately, from the point of view of philosophical foundations, I think the Bayesian position has won the day; if there is to be what we now call a coherent procedure, it has to be a Bayesian procedure.

2.2 Non-Bayesian and Non-Expected Utility Theories

Not long after the axiomatic foundations of expected utility given by von Neumann and Morgenstern and subjective expected utility in [Savage, 1954], the first objections started popping up. The first one is based on a violation of the independence axiom vNM2, and it is due to the French physicist Maurice Allais, who

was awarded the Nobel prize-equivalent for economics in 1988.⁵

Maurice Allais (1953) In economic theory, the name of Maurice Allais is famous mainly for the celebrated Allais' paradox ([Allais, 1953]). Allais imagined a series of two decision problems, each consisting in a choice between two options. The two decisions are depicted in the next table. In the first problem DM has to choose between lottery $1a$, that pays one million with certainty, and lottery $1b$, that pays one million with 89% probability, nothing with 1% probability, and five millions with 10% probability. Similarly, in the second problem DM has to choose between lotteries $2a$ and $2b$, with prizes according to the table below.

	89%	1%	10%
$1a$	1 million	1million	1 million
$1b$	1 million	0	5 millions
$2a$	0	1 million	1 million
$2b$	0	0	5 millions

Allais predicted that DM would choose lottery $1a$ in the first problem and $2b$ in the second, and subsequent experiments confirmed his intuition (see [Machina, 1987], [Oliver, 2003]). This behavioral pattern contradicts expected utility. Indeed, it is easy to see that the preference $1a \succ 1b$ in the first problem implies

$$u(\$1M) > 0.89u(\$1M) + 0.01u(\$0) + 0.1u(\$5M), \quad (2.3)$$

and consequently

$$0.11u(\$1M) - 0.1u(\$5M) > 0.01u(\$0).$$

On the other hand, preference $2a \prec 2b$ in the second problem implies that

$$0.9u(\$0) + 0.1u(\$5M) > 0.89u(\$0) + 0.11u(\$1M), \quad (2.4)$$

and so

$$0.01u(\$0) > 0.11u(\$1M) - 0.1u(\$5M),$$

which is in contradiction to what we got from the first decision problem. As anticipated, Allais' paradox violates the independence axiom vNM2. Consider the lotteries $L = \{\$1M, 1\}$, $L^2 = \{\$0, 1\}$, and $L^b = \{\$0M, 1/11; \$5M, 10/11\}$. Lotteries $1a$ and $1b$ can then be reformulated as

$$\begin{aligned} 1a &= L = 0.89L + 0.11L \\ 1b &= 0.89L + 0.11L^b \end{aligned} \quad ,$$

⁵The Nobel Memorial Prize in Economic Sciences is often referred to as the Nobel prize in economics, but to be precise it is not in the list of the five original Nobel prizes (physics, literature, chemistry, peace, and medicine), and it was established by Sweden's central bank much later, in 1968. Although, for the sake of brevity, sometimes we will refer to it as the Nobel prize in economics, it is worth remembering that there is no original Nobel prize for economics.

and similarly for lotteries $2a$ and $2b$:

$$\begin{aligned} 2a &= 0.89L^2 + 0.11L \\ 2b &= 0.89L^2 + 0.11L^b \end{aligned}$$

The preferences displayed by DM over the two problems are equivalent to

$$\begin{aligned} L &\succ 0.89L + 0.11L^b \\ 0.89L^2 + 0.11L &\prec 0.89L^2 + 0.11L^b \end{aligned}$$

It is a direct implication of the independence axiom vNM2 that $1a \succsim 1b$ if and only if $L \succsim L^b$. Hence, $1a \succ 1b$ implies $L \succ L^b$. But by the same reasoning, $2a \prec 2b$ implies $L \prec L^b$.

Few years after its establishment, the axiomatic solidity of expected utility began to creak under the weight of actual behavior.

Daniel Ellsberg (1961) Daniel Ellsberg is an economist and political activist. He is mainly known for two things. In 1971, while employed as U.S. military analyst, he decided to release the *Pentagon Papers*, a top-secret Pentagon document about the U.S. government’s decision making relative to the Vietnam War, to the *New York Times*. For this reason, he was charged of theft and conspiracy for a total sentence up to 115 years. In 1973, all charges were dismissed, and in 2006 he got the *Right Livelihood award* (also known as the *alternative Nobel prize*) for “offering practical and exemplary answers to the most urgent challenges facing us today” ([RLA, 1980]), also awarded, among others, to Mordechai Vanunu, Amy Goodman, Edward Snowden, and Gino Strada.

Daniel Ellsberg is also worldwide famous for Ellsberg’s paradoxes, which represent one of the hardest challenges to Savage’s axioms (see [Ellsberg, 1961]). Ellsberg’s single-urn paradox goes as follows. There is an urn containing 90 balls: 30 red balls, and the 60 remaining balls are either black or yellow, in unknown proportions. DM is faced with two decision problems. In each of them, a ball is drawn from the urn and DM wins \$100 if she guesses the color of the drawn ball. In the first decision problem, DM can choose between two options: betting on red r , and betting on black b . As for the second problem, DM still has two possible acts: betting on red or yellow ry , and betting on black or yellow by . The two situations are pictured in the following table. Which bet would you choose in the first problem? And in the second? Take your time!

	R	B	Y
r	\$100	0	0
b	0	\$100	0
ry	\$100	0	\$100
by	0	\$100	\$100

The prediction of Ellsberg, in line with the answers he observed when he proposed the experiment, is that the majority of the people would choose to bet

on r in the first decision problem and on by in the second. This behavior has later been tested and corroborated by experimental evidence (for a review see [Trautmann and van de Kuilen, 2016]).

Yet, the choices in Ellsberg's paradox are in contradiction with Savage's axioms, in particular with axiom P2, also known as the *sure-thing principle*. Let us define the act y in the obvious way, that is, y pays \$100 if a yellow ball is drawn and nothing otherwise. Then it is evident that acts ry and by of Ellsberg's second problem correspond to r_Y^y and b_Y^y , respectively. Hence, it follows from the sure-thing principle that: $r \succsim b$ if and only if $r_Y^y \succsim b_Y^y$. But we have seen that typically DM's preferences are such that $r \succ b$ and $r_Y^y \prec b_Y^y$. This pattern reflects an *aversion* to the ambiguity regarding the proportions in the urn. In both choices, DM picks the act that reduces that ambiguity as much as possible, as if she tried to secure herself against the worst possible urn compositions.

The behavior in Ellsberg's example clashes with the Bayesian paradigm and subjective expected utility theory (SEU henceforth). It is evident that there is no subjective probability distribution over the possible urn compositions that can justify such preferences. For that to be the case, we should find a probability P over the events R, B and Y such that

$$P(R)u(\$100) + P(B)u(0) + P(Y)u(0) > P(R)u(0) + P(B)u(\$100) + P(Y)u(0)$$

and

$$P(R)u(\$100) + P(B)u(0) + P(Y)u(\$100) < P(R)u(0) + P(B)u(\$100) + P(Y)u(\$100).$$

By simple calculation, the first inequality simplifies to $P(R) > P(B)$, whereas the second simplifies to $P(R) < P(B)$. Hence, there is no subjective probabilistic belief about the urn that can explain DM's behavior in terms of the maximization of subjective expected utility.

Ellsberg's intuition might bring us even deeper down the rabbit hole. If, for instance, we remember the two tenets of the Bayesian approach introduced earlier, and we want to maintain that the only sensible measure of subjective degrees of belief in some event is DM's willingness to bet, then we have to admit that in such cases DM's beliefs are not Bayesian, namely, DM's uncertainty cannot be expressed by a probability function. Moreover, many people seem to refrain from restoring probabilistic behavior even when one points to the inconsistency with Bayesian prescriptions. As reported by [Ellsberg, 1961],

The important finding is that, after rethinking all their "offending" decisions in the light of the axioms, a number of people who are not only sophisticated but reasonable decide that they wish to persist in their choices. This includes people who previously felt a "first-order commitment" to the axioms, many of them surprised and some dismayed to find that they wished, in these situations, to violate the

Sure-thing Principle. Since this group included L. J. Savage, when last tested by me (I have been reluctant to try him again), it seems to deserve respectful consideration.

This touches upon an important point in relation to Savage's axioms: their *normativity*. It may come as no surprise that axioms of subjective expected utility happen to be (maybe even regularly and consistently) violated in actual decision making. Decision makers are human, and humans are often fallible and irrational. After all, as a criterion for rationality, maximization of subjective expected utility does not claim any descriptive power. Instead, SEU axioms should be compelling from a normative point of view instead. But, on the contrary, many people who have been exposed to Ellsberg's experiment (including Savage himself), preferred to stick to their choices rather than conforming to the axioms.

What we have learned from Ellsberg's paradox is that sometimes DM has trouble forming Bayesian beliefs. This might depend, for example, on the quality of the information available. Whenever the information is clear and precise (in the case of an urn with known composition for instance), DM will most likely form Bayesian beliefs, and face a decision problem under risk. On the other hand, the more the available information becomes unreliable or unspecified, the less DM should be expected to boil the uncertainty down to risk. In these cases (such as Ellsberg's urn), DM is said to face a decision problem under unmeasurable uncertainty, or, in Ellsberg's terminology, *ambiguity*.

After [Ellsberg, 1961], ambiguity has occupied a central position in decision theory literature, and many frameworks have been proposed to encompass situations where DM's uncertainty is not reduced to risk. We will now suspend the chronological development for a moment, in order to stay on this track and elaborate more on decision making under ambiguity. We will have to jump back to the Allais' paradox in a few pages.

The idea of modeling DM's beliefs by means of multiple priors had already been introduced in decision theory at least since the work by [Dempster, 1967], [Levi, 1974], [Shafer, 1976], and [Gardenfors and Sahlin, 1982], but it is with the axiomatization given by [Gilboa and Schmeidler, 1989] that the connection between multiple priors and ambiguity became apparent.

Itzhak Gilboa and David Schmeidler (1989) [Gilboa and Schmeidler, 1989] prove a theorem that includes the possibility of DM entertaining non-Bayesian beliefs. In case of ambiguity, DM is represented as holding a (compact and convex) set of prior probability functions, and an act is evaluated by considering the minimal expected utility attained by that act, ranging over the set of possible priors. Accordingly, this approach (and the paper as well) is named *maxmin expected utility with non-unique prior*. Multiple-prior models have been very successful ever since, and they are nowadays the dominating paradigm for the representation of decisions under unmeasurable uncertainty.

The set-up in [Gilboa and Schmeidler, 1989] is similar to that in [Savage, 1954]. The main difference is that outcomes may be non-deterministic, usually called (*roulette*) *lotteries*. Roulette lotteries are defined exactly as the lotteries in [von Neumann and Morgenstern, 1944]. This amounts to saying that the outcome set is now $\Delta(X)$ instead of X , where $\Delta(X)$ is the set of all simple lotteries over X . Consequently, an act is defined as a function $f : S \rightarrow \Delta(X)$, where the state space S is endowed with an event algebra Σ , and acts are Σ -measurable finite step functions. Let us denote the set of all acts by \mathcal{F} and the set of constant acts by $\Delta(X)$ (obviously, $\Delta(X) \subset \mathcal{F}$). Everything else is inherited from Savage's framework, but it is to be remarked that the presence of lotteries as outcomes makes the outcome space convex. It follows that the set of acts \mathcal{F} is also convex. For $f, g \in \mathcal{F}$ and $\alpha \in (0, 1)$, the act $h = \alpha f + (1 - \alpha)g$ is defined pointwise by $h(s) = \alpha f(s) + (1 - \alpha)g(s)$ for all $s \in S$. Finally, the axioms on the preference relation $\succsim \subseteq \mathcal{F} \times \mathcal{F}$ are the following.

- GS1. *Total preorder*: \succsim is complete and transitive.
- GS2. *Certainty independence*: For $f, g \in \mathcal{F}$ and $L \in \Delta(X)$, and for $\alpha \in (0, 1)$, $f \succ g$ iff $\alpha f + (1 - \alpha)L \succ \alpha g + (1 - \alpha)L$.
- GS3. *Archimedean*: For $f, g, h \in \mathcal{F}$, if $f \succ g \succ h$, then $\exists \alpha, \beta \in (0, 1)$ such that $\alpha f + (1 - \alpha)h \succ g \succ \beta f + (1 - \beta)h$.
- GS4. *Monotonicity*: For $f, g \in \mathcal{F}$, if $\forall s \in S f(s) \succsim g(s)$ then $f \succsim g$.
- GS5. *Ambiguity aversion*: For $f, g \in \mathcal{F}$ and $\alpha \in (0, 1)$, $f \sim g$ implies $\alpha f + (1 - \alpha)g \succsim f$.
- GS6. *Nontriviality*: $\exists f, g \in \mathcal{F}$ such that $f \succ g$.

The representation result in [Gilboa and Schmeidler, 1989] states that the following are equivalent:

1. \succsim satisfies axioms GS1-GS6.
2. There exists a non-constant and unique (up to positive affine transformations) function $u : X \rightarrow \mathbb{R}$ and a compact and convex set $\Gamma \subseteq \Delta(\Sigma)$ of finitely additive probability measures such that, for all $f, g \in \mathcal{F}$,

$$f \succsim g \text{ iff } \min_{P \in \Gamma} \int_S \bar{u}(f(s)) dP(s) \geq \min_{P \in \Gamma} \int_S \bar{u}(g(s)) dP(s),$$

where $\bar{u}(L) = \sum_{x \in \text{supp}L} u(x)L(x)$, as in von Neumann-Morgenstern's representation theorem.

GS2, together with GS5, is the key axiom. Certainty independence restricts the axiom of independence of [Anscombe and Aumann, 1963] to convex mixtures with constant acts only. Loosely speaking, once the sure-thing principle has been weakened to certainty independence, DM's typical behavior in Ellsberg's example is no longer in contradiction with the axioms.⁶ To show this in the present setup, let us define the auxiliary acts r^* , b^* , y^* and f^0 such that r^* pays \$200 if a red ball is drawn and nothing otherwise, b^* pays \$200 if a black ball is drawn and nothing otherwise, y^* pays \$200 if a yellow ball is drawn and nothing otherwise, and act f^0 pays nothing in any case. We can then express acts r and b in Ellsberg's paradox as $r = 0.5r^* + 0.5f^0$ and $b = 0.5b^* + 0.5f^0$, and acts ry and by as $ry = 0.5r^* + 0.5y^*$ and $by = 0.5b^* + 0.5y^*$. If we allowed act h of axiom GS2 to be any act, then we would have that: $r^* \succ b^* \Leftrightarrow r \succ b \Leftrightarrow ry \succ by$. Ellsberg's paradox would not be solved. But since certainty independence just admits mixtures with constant acts, we have instead: $r^* \succ b^* \Leftrightarrow r \succ b \not\Leftrightarrow ry \succ by$. A way out of the impasse is found. We can consequently think of DM as maximinimizing expected utility over the obvious set of priors

$$\Gamma = \{P \in \Delta(\{R, B, Y\}) : P(R) = \frac{1}{3}, P(B) = \alpha, P(Y) = 1 - P(R) - \alpha, \text{ for } \alpha \in [0, \frac{2}{3}]\}.$$

This amounts to computing, for any probability measure in Γ , the expected utility of each act, and then choosing the act that guarantees the highest minimal expected utility. The minimal expected utility of r is then \$33.33 and the minimal expected utility of b is 0, whereas the minimal expected utility of ry is \$33.33 and the minimal expected utility of by is \$66.66: DM will prefer r over b and by over ry .

As already mentioned, experimental evidence suggests that the behavior predicted by Ellsberg is rather regular and systematic, and the axiomatization of [Gilboa and Schmeidler, 1989] is hence an unquestionable contribution from a descriptive point of view. But is that to be viewed as descriptive theory only? In our opinion, the answer is negative. We agree with [Gilboa and Marinacci, 2013] in that

DMs may sometimes admit that they do not know what the probabilities they face are. Being able to admit ignorance is not a mistake. It is, we claim, more rational than to pretend that one knows what cannot be known.

[...] When central bank executives consider monetary policies, and when leaders of a country make decisions about military actions, they will not make a mistake if they do not form Bayesian probabilities.

⁶The weakening represented by certainty independence would be more evident had we introduced the Anscombe-Aumann version of SEU (see [Anscombe and Aumann, 1963]). We refer the interested reader directly to the original paper for more details.

On the contrary, they will be well advised to take into account those uncertainties that cannot be quantified.

Therefore, we maintain, multiple-prior models are compatible with *normative* interpretations too.

As a final curiosity, the 2016 annual report on global catastrophic risks developed by the joint work of the Global Challenges Foundations and the Future of Humanity Institute at Oxford listed the twelve biggest threats that might lead to the end of the world (see [GPP, 2016]). The following table summarizes their conclusions.

possible cause	probability	possible cause	probability
climate change	0,01%	super-volcanoes	0,00003%
nuclear war	0,005%	asteroids	0,00013%
natural pandemics	0,0001%	biotechnology	0,01%
economic collapse	unknown	eco-catastroph	unknown
engineering pandemic	0,01%	global misrule	unknown
artificial intelligence	0-10%	unknown risks	0,1%

This report furnishes a clear example of the two types of uncertainty just discussed. While, boldly enough, the research attaches precise probabilistic risk to some of the possible causes, like super-volcanoes and nuclear war, for others it can only step back from Bayesianism and admit ignorance.

Daniel Kahneman and Amos Tversky (1979) Ten years before the work by [Gilboa and Schmeidler, 1989], an article that addressed the issues raised by Allais' paradox and similar experiments appeared in *Econometrica* under the title of *Prospect Theory: An Analysis of Decision under Risk*. The authors were two Israeli psychologists, Daniel Kahneman and Amos Tversky, and their paper, [Kahneman and Tversky, 1979], became the most cited paper ever published in *Econometrica*. Daniel Kahneman was also awarded the Nobel prize-equivalent for economics in 2002.

The flavor of Kahneman and Tversky's Prospect theory is rather descriptive than normative. The theory is originated from

several classes of choice problems in which preferences systematically violate the axioms of expected utility theory. In the light of these observations we argue that utility theory, as it is commonly interpreted and applied, is not an adequate descriptive model and we propose an alternative account of choice under risk. ([Kahneman and Tversky, 1979])

A *prospect* $X_i = (x_1, p_1; \dots; x_n, p_n)$ is a lottery that yields monetary outcome x_j with probability p_j , for $1 \leq j \leq n$ and $\sum_{j=1}^n p_j = 1$. [Kahneman and Tversky, 1979] take into consideration only choices between binary prospects, i.e., prospects with at most two nonzero outcomes. Table 2.1 (borrowed from [Loomes and Sugden, 1982])

Kahneman and Tversky problem no.	Prospects offered†	Modal preference	Percentage of subjects with modal preference	Characterisation of modal preference
1	$X_1 = (2,500, 0.33; 2,400, 0.66)$	$X_1 \prec X_2$	82*	Risk averse
2	$X_2 = (2,400, 1.00)$ $X_3 = (2,500, 0.33)$ $X_4 = (2,400, 0.34)$	$X_3 \succ X_4$	83*	Not clear
3	$X_5 = (4,000, 0.80)$ $X_6 = (3,000, 1.00)$	$X_5 \prec X_6$	80*	Risk averse
3'	$X_7 = (-4,000, 0.80)$ $X_8 = (-3,000, 1.00)$	$X_7 \succ X_8$	92*	Risk loving
4	$X_9 = (4,000, 0.20)$ $X_{10} = (3,000, 0.25)$	$X_9 \succ X_{10}$	65*	Not clear
4'	$X_{11} = (-4,000, 0.20)$ $X_{12} = (-3,000, 0.25)$	$X_{11} \prec X_{12}$	58	Not clear
7	$X_{13} = (6,000, 0.45)$ $X_{14} = (3,000, 0.90)$	$X_{13} \prec X_{14}$	86*	Risk averse
8	$X_{15} = (6,000, 0.001)$ $X_{16} = (3,000, 0.002)$	$X_{15} \succ X_{16}$	73*	Risk loving
10	$X_{17} = (X_5, 0.25)$ $X_{18} = (X_6, 0.25)$	$X_{17} \prec X_{18}$	78*	Risk averse
14	$X_{19} = (5,000, 0.001)$ $X_{20} = (5, 1.000)$	$X_{19} \succ X_{20}$	72*	Risk loving
14'	$X_{21} = (-5,000, 0.001)$ $X_{22} = (-5, 1.000)$	$X_{21} \prec X_{22}$	83*	Risk averse

* Statistically significant at the 0.01 level.

† Consequences are increments or decrements of wealth, measured in Israeli pounds.

Table 2.1: Experimental findings of [Kahneman and Tversky, 1979].

is a selection of some of the experimental results in [Kahneman and Tversky, 1979], where prospect (2.500, 0.33) is an abbreviation for (2.500, 0.33; 0, 0.67). The principal findings from these experiments are essentially three:

1. *Allais' paradox* and the *certainty effect*: corresponding to the combination of $X_1 \prec X_2$ and $X_3 \succ X_4$, the combination of $X_5 \prec X_6$ and $X_9 \succ X_{10}$, and the combination of $X_{13} \prec X_{14}$ and $X_{15} \succ X_{16}$. The *reverse certainty effect* is also present, in the combination of $X_7 \succ X_8$ and $X_{11} \prec X_{12}$.
2. The *isolation effect* in two-stage problems: in the combination $X_9 \succ X_{10}$ and $X_{17} \prec X_{18}$.
3. The *reflection effect*: in the combination of $X_5 \prec X_6$ and $X_7 \succ X_8$, in the combination of $X_9 \succ X_{10}$ and $X_{11} \prec X_{12}$, and in the combination of $X_{19} \succ X_{20}$ and $X_{21} \prec X_{22}$.

While the reflection effect is not an explicit violation of classic expected utility theory, Allais' paradox and the certainty effect as well as the isolation effect openly violate expected utility axioms (as we have already seen for Allais' paradox).⁷

⁷For reasons of succinctness, we will not talk at more length about these effects here. The certainty effect is exemplified by Allais' paradox, which has already been discussed above. For a more detailed account of the other effects we refer the reader to [Kahneman and Tversky, 1979].

The new theory proposed by Kahneman and Tversky to accommodate all these phenomena is built on two central pivots: a value function $v : X \rightarrow \mathbb{R}$ and a weighting function $\pi : [0, 1] \rightarrow [0, 1]$. As for the *value function* v , the carrier of value $v(x)$ is the change in wealth at outcome x , rather than the final amount of money per se. Then, the indifferent point with respect to changes in wealth is $x = 0$, where DM does not incur any gains or losses, such that $v(0) = 0$. This is called the *reference point*, and experiments suggest that the value function is concave in the domain of gains ($v''(x) < 0$ for $x > 0$), and convex in the domain of losses ($v''(x) > 0$ for $x < 0$).⁸ Moreover, losses loom larger than gains, that is, v is steeper below the reference point than above ($v'(x) < -v'(-x)$ for $x > 0$). This would explain why many people find symmetric lotteries $(x, 0.50; -x, 0.50)$ rather unattractive.

As to the *weighting function*, it is assumed that DM's probabilities p are distorted by a function π . Although it is still required that $\pi(0) = 0$ and $\pi(1) = 1$, the resulting decision weights $\pi(p)$ do not obey the axioms of probability theory. In particular, given the concavity of v , problems like 8 in Table 2.1 suggest that, for small values p , the function π is *subadditive*, namely, $\pi(\alpha p) > \alpha\pi(p)$ for $\alpha \in (0, 1)$. Moreover, because of certainty effects it seems that DM overestimates low probabilities, so that $\pi(p) > p$ for low p . Nevertheless, the weighting function is also assumed to satisfy *subcertainty*: for all $p \in (0, 1)$, $\pi(p) + \pi(1 - p) < 1$.

Finally, an important distinction is the one between *regular* prospects and *non-regular* prospects. Non-regular prospects are prospects of the form $(x, p; y, q)$ with $p + q = 1$ and either $x, y > 0$ or $x, y < 0$. Prospects that are not non-regular are regular. The value of a prospect is given by a function V , depending on the type of the prospect:

- for regular prospects:

$$V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y);$$

- for non-regular prospects: for either $x > y > 0$ or $x < y < 0$,

$$V(x, p; y, q) = v(y) + \pi(p)(v(x) - v(y)).$$

If we go back to Allais' paradox for a moment, we can have an idea of how it can be accommodated within Prospect theory. By rewriting equations (2.3) and (2.4) according to the appropriate V we get

$$v(\$1M) > \pi(0.89)v(\$1M) + \pi(0.01)v(\$0) + \pi(0.1)v(\$5M),$$

and

$$\pi(0.9)v(\$0) + \pi(0.1)v(\$5M) > \pi(0.89)v(\$0) + \pi(0.11)v(\$1M).$$

⁸This assumption produces the celebrated S-shaped utility function of Prospect theory.

The first simplifies to

$$v(\$1M) - \pi(0.89)v(\$1M) > \pi(0.1)v(\$5M),$$

while the second reduces to

$$\pi(0.1)v(\$5M) > \pi(0.11)v(\$1M).$$

By combining the two we have

$$v(\$1M) - \pi(0.89)v(\$1M) > \pi(0.11)v(\$1M),$$

and consequently

$$1 > \pi(0.89) + \pi(0.11).$$

In words, if DM has a function π that distorts the probabilities associated to different outcomes, and that satisfies subcertainty, then Allais' paradox can be solved. The assumption that $\pi(p) > p$ for low p , together with $\pi(p) + \pi(1-p) < 1$, reflects DM's preference switch in Allais' examples and justify the behavior associated with the certainty effect: 1% difference in probability from 100% to 99% is much more significant than from 11% to 10%.

As it should be apparent by now, Prospect theory has no normative claims: there is no compelling reason why DM *should* turn known probabilities into non-probabilistic decision weights for example. The goal of the theory is purely descriptive: it aims at formalizing a suitable mathematical framework to encompass actual observed behavior, without providing any reason or explanation, apart from experimental observations, for the decision rules introduced.

Graham Loomes and Robert Sugden (1982) An article written by two British economists, Graham Loomes and Robert Sugden, entitled *Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty* appeared just a few years after [Kahneman and Tversky, 1979], and, in our opinion, would have deserved as much attention as the paper by Kahneman and Tversky. The idea of [Loomes and Sugden, 1982] is straightforward. They realized that there is a reason, or a principle, that could explain and connect all the empirical findings that originated Prospect theory: the *regret*. Without performing any new experiment, they simply borrowed the behavioral evidence from Kahneman and Tversky and showed how a regret-based theory of choice can accommodate DM's observed behavior. Their point is that DM's behavior is better explained as aiming to minimize expected regret rather than to maximize expected utility. Roughly speaking, the regret is the loss that DM incurs when the action she chooses turns out not to be optimal. Minimizing this potential loss seems a perfectly legitimate goal for DM's decision making, and it can offer powerful insights also from a descriptive and explanatory point of view. Regret-based reasoning will play a crucial role in next chapters, so let's zoom in on their ideas in more detail.

Let's assume for simplicity that there is a finite number n of states of the world s_1, \dots, s_n . According to expected utility, an act, or action, f is preferred to action g if f yields higher expected utility than g ,

$$f \succsim g \text{ iff } \sum_{i=1}^n u(f(s_i))P(s_i) \geq \sum_{i=1}^n u(g(s_i))P(s_i).$$

The last formula is of course equivalent to

$$f \succsim g \text{ iff } \sum_{i=1}^n P(s_i)(u(f(s_i)) - u(g(s_i))) \geq 0. \quad (2.5)$$

According to regret theory instead,

$$f \succsim g \text{ iff } \sum_{i=1}^n P(s_i)Q(u(f(s_i)) - u(g(s_i))) \geq 0, \quad (2.6)$$

where Q is a strictly increasing function of the difference in utilities $u(f(s_i)) - u(g(s_i))$, and represents the regret-rejoice of f relative to g at s_i . Besides, the regret-rejoice function Q is supposed to have the following symmetry property: $Q(-x) = -Q(x)$. Consequently, $Q(0) = 0$. When Q is linear, it holds that

$$Q(u(f(s_i)) - u(g(s_i))) = Q(u(f(s_i))) - Q(u(g(s_i))),$$

and if Q is linear and strictly increasing, it is a positive affine transformation. Furthermore, we know from Savage's theorem that u is unique up to positive affine transformations. Hence, when Q is a linear function, regret theory and expected utility theory are indistinguishable. The assumption of [Loomes and Sugden, 1982] is instead that Q is convex on \mathbb{R}_+ (and, therefore, concave on \mathbb{R}_-).

To see how regret theory is able to accommodate the behavioral findings of Prospect theory, consider Allais' paradox again. Written in the form of prospects, the typical preferences in Allais' two problems are:

1. $1a = (\$1M, 1) \succ (\$1M, 0.89; \$5M, 0.1) = 1b$;
2. $2a = (\$1M, 0.11) \prec (\$5M, 0.1) = 2b$.

The state space for the first problem is generated by the possible combinations of outcomes from prospects $1a$ and $1b$, as described in the next table.

	$1a$	$1b$	$P(s_i)$
s_1	$\$1M$	$\$5M$	0.1
s_2	$\$1M$	$\$1M$	0.89
s_3	$\$1M$	$\$0$	0.01

According to regret theory, DM's preference $1a \succ 1b$ is equivalent to

$$0.1Q(u(\$1M) - u(\$5M)) + 0.89Q(u(\$1M) - u(\$1M)) + 0.01Q(u(\$1M) - u(\$0)) > 0.$$

By simple computation, the previous inequality is in turn equivalent to

$$Q(u(\$5M) - u(\$1M)) < 0.1Q(u(\$1M) - u(\$0)). \tag{2.7}$$

Intuitively, this means that the regret from winning one million instead of five millions is more than ten times smaller compared to the regret from getting nothing instead of one million. When we come to consider Allais' second problem, the state space is as shown in the following table.

	$2a$	$2b$	$P(s_i)$
s_1	$\$1M$	$\$5M$	0.011
s_2	$\$1M$	$\$0$	0.099
s_3	$\$0$	$\$5M$	0.089
s_4	$\$0$	$\$0$	0.801

The typical preference $2a \prec 2b$ observed in the second problem would then correspond to

$$0.011Q(u(\$1M) - u(\$5M)) + 0.099Q(u(\$1M) - u(\$0)) + 0.089Q(u(\$0) - u(\$5M)) < 0,$$

which simplifies to

$$Q(u(\$5M) - u(\$1M)) > 9Q(u(\$1M) - u(\$0)) - \frac{89}{11}Q(u(\$5M) - u(\$0)). \tag{2.8}$$

Putting (2.7) and (2.8) together, we obtain

$$\frac{10}{11}Q(u(\$5M) - u(\$0)) > Q(u(\$1M) - u(\$0)).$$

Therefore, DM's regret from getting nothing instead of one million is now smaller than the regret from getting nothing instead of five millions, and rightfully so. Evidently, expected utility theory would fail in accommodating DM's behavior. As we have seen, expected utility theory amounts to the linearity of Q . If Q is linear, equation (2.7) reduces to

$$Q(5) < \frac{11}{10}Q(1) - \frac{1}{10}Q(0),$$

while equation (2.8) can be reduced to

$$Q(5) > \frac{11}{10}Q(1) - \frac{1}{10}Q(0).$$

In full generality, we can express Allais' paradox and the certainty effect as the choice between $a = (\$x, p + \alpha)$ and $b = (\$x, \alpha; \$y, q)$, where $0 < x < y$, $0 < q < p \leq 1$, and $\alpha \leq (1 - p)$, such that:

1. $a \succ b$, when $\alpha = (1 - p)$;
2. $a \prec b$, when $\alpha = 0$.

The general state space is generated in the obvious way, as specified by the following table.

	a	b	$P(s_i)$
s_1	$\$x$	$\$x$	$\alpha(\alpha + p)$
s_2	$\$x$	$\$y$	$q(\alpha + p)$
s_3	$\$x$	$\$0$	$(1 - \alpha - q)(\alpha + p)$
s_4	$\$0$	$\$x$	$\alpha(1 - \alpha - p)$
s_5	$\$0$	$\$y$	$q(1 - \alpha - p)$
s_6	$\$0$	$\$0$	$(1 - \alpha - q)(1 - \alpha - p)$

It is clear that if we set $p = 0.11$, $q = 0.1$, $x = 1M$ and $y = 5M$, we are back to Allais' paradox. For any given p, q, x, y , if we denote by $\bar{\alpha}$ the value such that $a \sim b$ when $\alpha = \bar{\alpha}$, then we will observe $a \succ b$ for $\alpha > \bar{\alpha}$, and $a \prec b$ for $\alpha < \bar{\alpha}$.

It is important to remark that the specification of the options in the form of prospects underlies that the lotteries are independent and come from different drawings. Independence of lotteries is fundamental for regret to give a different prediction than expected utility in Allais' paradox. Let us suppose that the two lotteries $1a$ and $1b$, as well as the two lotteries $2a$ and $2b$, are not independent, but rather depend on the same drawing of one ball from an urn with 100 numbered balls. The prizes of the four lotteries are listed in the following table accordingly.

	1	89	90	91	100
$1a$	1 million	1million	1 million	1 million	
$1b$	1 million		0	5 millions	
$2a$	0		1 million	1 million	
$2b$	0		0	5 millions	

When DM faces this version of the paradox, regret theory would also predict that $1a \succ 1b$ if and only if $2a \succ 2b$. If the lotteries are not independent, then the state spaces would differ from those specified above, and the utilities at state $1 - 89$ would always cancel out in the summation. Whenever two actions achieve the same outcome at a state, their utilities automatically cancel out according to regret theory: $Q(u(f(s)) - u(g(s))) = 0$ whenever $f(s) = g(s)$. For this reason, regret theory satisfies the sure-thing principle:

$$f_{E^c}^h \succsim g_{E^c}^h \text{ iff } f_{E^c}^{h'} \succsim g_{E^c}^{h'}.$$

It does not satisfy the independence axiom vNM2, though. Indeed, the prospect formulation of Allais' paradox is equivalent to the choice between lotteries that we presented in the section on Maurice Allais:

1. $1a = L = (\$1M, 1) \succ (\$1M, 0.89; \$5M, 0.1) = 0.89L + 0.11L^b = 1b$;

$$2. \quad 2a = 0.89L^2 + 0.11L = (\$1M, 0.11) \prec (\$5M, 0.1) = 0.89L^2 + 0.11L^b = 2b.$$

Interestingly, [Birnbbaum, 2008] reported that the typical behavior in Allais’ paradox is more systematic when the lotteries are presented as independent drawings.

Here, we have decided to focus only on Allais’ paradox, as the working example, in order to illustrate how regret theory can explain Kahneman and Tversky’s experimental observations, but all the other empirical findings of Prospect theory discussed earlier can be accommodated by regret theory as well. However, when compared to Prospect theory, regret theory has the advantage of exhibiting a legitimate decision principle, the minimization of (some transformation of) regret, that might lead DM to behave at odds with expected utility maximization.

Before concluding, it is worth mentioning another main deviation from expected utility theory. Consider the next table: a ball will be drawn from an urn containing 90 numbered balls.

	1	30	31	60	61	90
<i>a</i>		\$0		\$100		\$200
<i>b</i>		\$100		\$200		\$0
<i>c</i>		\$200		\$0		\$100

According to expected utility, DM should be indifferent between acts *a* and *b*, since they have the same expected utility:

$$\frac{1}{3}u(\$0) + \frac{1}{3}u(\$100) + \frac{1}{3}u(\$200) = \frac{1}{3}u(\$100) + \frac{1}{3}u(\$200) + \frac{1}{3}u(\$0).$$

This is precisely the prescription of the *equivalence axiom*: acts that induce the same probabilities over outcomes are equivalent. Regret theory, however, does not satisfy the equivalence axiom. Indeed, it follows from the convexity of Q on \mathbb{R}_+ that

$$Q(u(\$200) - u(\$0)) > Q(u(\$200) - u(\$100)) + Q(u(\$100) - u(\$0)),$$

therefore

$$\frac{1}{3}Q(u(\$0) - u(\$100)) + \frac{1}{3}Q(u(\$100) - u(\$200)) + \frac{1}{3}Q(u(\$200) - u(\$0)) > 0,$$

and consequently $a \succ b$, even if *a* and *b* induce exactly the same probabilities over the outcomes. Moreover, by the same reasoning, regret theory predicts $b \succ c$ and $c \succ a$, which constitutes a preference cycle. This pattern violates transitivity and, hence, the basic total preorder axiom.⁹ It should then be no surprise for the reader who remembers the quote from [Bather, 1996] about Savage that regret theory also violates the *independence of irrelevant alternatives* (IIA): if an option

⁹This axiom is supposed to be so essential to rational choice that some authors (e.g., [Mas-Colell et al., 1995]) name it the “rationality axiom”.

a is chosen from the set of options A , then a should be chosen also from A' , for any $A' \subset A$ such that $a \in A'$. The violation of IIA manifests one of the most fundamental features of regret: its *context dependency*.

We want to conclude this paragraph with a brief discussion on the status of regret theory. While the contribution from a descriptive point of view is significant and indisputable, there is debate whether a normative reading of the theory is also viable. [Loomes and Sugden, 1982] explicitly advocate a normative interpretation, but [Bleichrodt and Wakker, 2015] for instance disagree on this point. In [Loomes and Sugden, 1982]’s words:

Thus we believe that regret theory does more than predict certain systematic violations of conventional expected utility theory: it indicates that such behaviour is not, in any meaningful sense of the word, irrational.

[...] we shall challenge the idea that the conventional axioms constitute the only acceptable basis for rational choice under uncertainty. We shall argue that it is no less rational to act in accordance with regret theory, and that conventional expected utility theory therefore represents an unnecessarily restrictive notion of rationality.

[...] it seems to us that psychological experiences of regret and rejoicing cannot properly be described in terms of the concept of rationality: a choice may be rational or irrational, but an experience is just an experience. As far as the second assumption is concerned, if an individual does experience such feelings, we cannot see how he can be deemed irrational for consistently taking those feelings into account.

[Bleichrodt and Wakker, 2015] instead, in contrast to this position, argue that

Taking any emotion as rational just because it exists is too permissive and applies Hume’s adage ‘reason is, and ought only to be the slave of the passions’ too leniently.

To position ourselves in the debate, we shall say that we are not at all hostile to the normative interpretation. On the contrary, we largely support it. The motivation behind our stance is different from [Loomes and Sugden, 1982] though, in that we rely on a theoretical, rather than psychological, argument. We think of the matter at hand as relating to a more general issue about the status of context dependency in rational choice. In our opinion, there is nothing wrong or disappointing with ascribing a normative status to context dependent reasoning. We suggest that there is no compelling reason to a priori deem as irrational a person whose evaluation of an option may radically change depending on which other options are concurrently available. We would even dare to claim that it would not be rational to do otherwise. Apart from the obvious descriptive merits, regret is one of the most successful instances of this context dependent approach to rational choice.

As anticipated when quoting Chernoff earlier, regret had already been introduced in decision theory by [Savage, 1951], but for a long time it has been out-classed by expected utility theory. As a final historical remark, 1982 was a great year of Renaissance for regret minimization. Other than [Loomes and Sugden, 1982], other two important papers, by [Bell, 1982] and [Fishburn, 1982], appeared that year, and breathed new life into regret theory.

Jörg Stoye (2011) For the sake of completeness, we have to mention that different representation theorems for regret-based choices have recently been presented. We are referring in particular to [Hayashi, 2008] and [Stoye, 2011].

We have already stressed that a main feature of regret is its context dependency, that is the reason of the violation of IIA. As a consequence, it is impossible for regret to fit into representations in terms of a binary preference relation \succsim . For instance, the last example above displays intransitive preferences $a \succ b, b \succ c$ and $c \succ a$. But to better explain this pattern, we shall explicitly notice that regret theory predicts that a is chosen over b when the available options are a and b only; similarly, b is preferred to c when the choice is between b and c only, and c is chosen over a when the choice is only between a and c . On the other hand, when all three alternatives are simultaneously available, DM is indifferent between them, $a \sim b \sim c$. This context dependent behavior cannot be represented by a context independent preference relation \succsim . For this reason, both [Hayashi, 2008] and [Stoye, 2011] have to state their representation theorems in terms of a *choice correspondence* $C : 2^{\mathcal{F}} \rightarrow 2^{\mathcal{F}}$. Such a choice correspondence can encode behavior that is dependent on the given menu of choices at hand. The previous example may thus be described by means of the following choice correspondence.

$$\begin{aligned} C(\{a, b\}) &= \{a\} & C(\{b, c\}) &= \{b\} \\ C(\{a, c\}) &= \{c\} & C(\{a, b, c\}) &= \{a, b, c\} \end{aligned}$$

Notice that these choices are still in violation of IIA: since $a \in C(\{a, b, c\})$ and $a \in \{a, c\} \subset \{a, b, c\}$, then by IIA $a \in C(\{a, c\})$, which is not the case.

The representation theorems for regret minimization presented in this section follow the formulation in [Stoye, 2011], but the interested reader should also see [Hayashi, 2008], [Milnor, 1954], and [Sugden, 1993]. The set-up in [Stoye, 2011] is the same as in [Gilboa and Schmeidler, 1989], but DM is now offered finite and nonempty menus of actions $M \subseteq \mathcal{F}$. DM's choice from menu M is represented by the choice correspondence $C(M) \subseteq M$. Finally, the following notation will be useful. Given menu M , action f and $\alpha \in [0, 1]$, we denote by $\alpha M + (1 - \alpha)f$ the menu generated by the convex combination of all actions in M with action f , i.e., $\alpha M + (1 - \alpha)f := \{\alpha g + (1 - \alpha)f : g \in M\}$. Moreover, let us define the *conditional* choice correspondence $C_s(M) \subseteq M$ at state $s \in S$ such that

$$f \in C_s(M) \text{ iff } f(s) \in C(\{g(s) : g \in M\}).$$

Notice that $f(s)$ and $g(s)$ are (roulette) lotteries, namely, $f(s), g(s) \in \Delta(X)$. Hence, the correspondence C_s assumes that DM's choice after state s is revealed corresponds to DM's choice from the constant acts $\{g(s) : g \in M\}$. We then say that an act f is *strictly potentially optimal* in M if there exists $s \in S$ such that $C_s(M) = \{f\}$.

In [Stoye, 2011], three possible versions of regret minimization are taken into account. Regret minimization with *no priors* corresponds to the proposal by [Savage, 1951]. The intuition is that the uncertainty in a given situation is represented directly and exclusively by the state space S . The other two versions of regret minimization instead admit a representation of the uncertainty in terms of a set of priors Γ over the space (S, Σ) . The difference between these two is that the set Γ may be endogenous or exogenous. The case of regret minimization with *endogenous priors* is parallel to the approach taken by [Gilboa and Schmeidler, 1989], where both DM's subjective utility u and subjective beliefs Γ are derived from the axioms. Conversely, the idea behind regret minimization with *exogenous priors* is that the ambiguity inherent in a decision situation is triggered by the environment in an objective way, and it is specified by an exogenously given set of priors Γ . In this latter case, there is no subjective representation of DM's beliefs, since Γ is already given by the environment from the beginning.

The axioms on the choice correspondence C used by Stoye for the representation of regret minimization with no priors are the following.

- S1. *Nontriviality*: For some $M \subseteq \mathcal{F}$, $C(M) \subset M$.
- S2. *Monotonicity*: If $f \in M$, $g \in C(M)$ and $f \in C_s(\{f, g\})$ for all $s \in S$, then $f \in C(M)$.
- S3. *Independence*: $C(\alpha M + (1 - \alpha)f) = \alpha C(M) + (1 - \alpha)f$.
- S4. *IIA for constant acts*: Let M and N consist of constant acts, then
- $$C(M \cup N) \cap M \in \{C(M), \emptyset\}.$$
- S5. *Independence of never strictly optimal alternatives*: Let M and N be such that $C_s(M \cup N) \cap M \neq \emptyset$ for all s , then
- $$C(M \cup N) \cap M \in \{C(M), \emptyset\}.$$
- S6. *Mixture continuity*: Fix any menu M and acts $g \in M$, $h \in \mathcal{F}$, and $f \notin M$ such that $C(M \cup \{f\}) = \{f\}$. Then there exists $\alpha \in (0, 1)$ such that

$$C(M \cup \{\alpha f + (1 - \alpha)h\}) = \{\alpha f + (1 - \alpha)h\}$$

and

$$\alpha g + (1 - \alpha)h \notin C(M \cup \{f, \alpha g + (1 - \alpha)h\}).$$

S7. *Ambiguity aversion*: $C(M)$ is the intersection of M with a convex set. That is, for any acts f, g , any $\alpha \in [0, 1]$, and any menu $M \supseteq \{f, g, \alpha f + (1 - \alpha)g\}$,

$$\text{if } f, g \in C(M) \text{ then } \alpha f + (1 - \alpha)g \in C(M).$$

S8. *Symmetry*: For any menu M and disjoint events $E_1, E_2 \in \Sigma \setminus \emptyset$ such that any $f \in M$ is constant on E_1 as well as on E_2 , let f' be defined by

$$f'(s) = \begin{cases} f(s)|_{s \in E_2}, & s \in E_1 \\ f(s)|_{s \in E_1}, & s \in E_2 \\ f(s) & \text{otherwise.} \end{cases}$$

Then the function $(\cdot)'$: $2^{\mathcal{F}} \rightarrow 2^{\mathcal{F}}$, that maps any $N \subseteq M$ onto $N' = \{f' : f \in N\}$, is such that

$$C(M') = (C(M))'.$$

The first representation theorem proven in [Stoye, 2011] states that, whenever Σ contains at least three events, the choice correspondence C fulfills axioms S1-S8 if and only if it can be represented as

$$C(M) = \operatorname{argmin}_{f \in M} \max_{s \in S} \{ \max_{g \in M} \bar{u}(g(s)) - \bar{u}(f(s)) \},$$

where the function \bar{u} is non-constant and unique up to positive affine transformations.

The prior-less feature of this result emerges from the second maximization over the state space S directly. We will see immediately that in case of regret minimization with endogenous and exogenous priors the formula obtained in the theorem includes a maximization over a set of priors Γ rather than over states. In order to get there, we need to define the notion of *state-independent outcome distribution*. If the set of outcomes $\{L \in \Delta(X) : \exists f \in M \text{ s.t. } f(s) = L\}$ is constant across all states s , then we say that menu M has state-independent outcome distribution. A menu consisting of constant acts necessarily has state-independent outcome distribution, but this property can be satisfied also by menus without any constant act. For example, for some $E \in \Sigma \setminus \emptyset$, the menu $M = \{f, g\}$ such that

$$\begin{aligned} f(s) &= L & \text{if } s \in E \\ f(s) &= L' & \text{if } s \in \Sigma \setminus E \\ g(s) &= L' & \text{if } s \in E \\ g(s) &= L & \text{if } s \in \Sigma \setminus E \end{aligned}$$

has state dependent outcome distribution without containing any constant act. Let us then introduce the next axiom:

S9. *C-betweenness when outcome distributions are state-independent:* For $f \in \mathcal{F}$, $L \in \Delta(X)$, $\alpha \in (0, 1)$, and menu $M \supseteq \{L, f, \alpha f + (1 - \alpha)L\}$ with state-independent outcome distribution,

$$\text{if } L \notin C(M) \text{ and } f \notin C(M), \text{ then } \alpha f + (1 - \alpha)L \notin C(M).$$

This axiom is dual to axiom S7. While ambiguity aversion guarantees that if two acts are ranked indifferent then neither of them can be chosen over any mixture of the two, c-betweenness ensures the opposite: if the menu has state-independent outcome distribution, then the mixture is not chosen if neither of them is. Axioms S1-S7 and S9 are satisfied if and only if the choice correspondence C can be represented as

$$C(M) = \operatorname{argmin}_{f \in M} \max_{P \in \Gamma} \int_S (\max_{g \in M} \bar{u}(g(s)) - \bar{u}(f(s))) dP(s),$$

where Γ and \bar{u} are as in [Gilboa and Schmeidler, 1989], i.e., Γ is compact, convex and unique, and \bar{u} is non-constant and unique up to positive affine transformations.

This result expresses the representation of regret minimization with endogenous priors, in that the set of priors Γ is uniquely defined by the axioms, namely, it is derived from DM's choices, rather than specified a priori by the environment.

As opposed to a subjective set of priors, there is the case of exogenous priors. Before stating the relevant axioms, we have to introduce the binary relation \triangleright_C for choice correspondence C . Specifically, we define $f \triangleright_C g$ if and only if, for all $\alpha \in (0, 1)$, $L \in \Delta(X)$, and $M \supseteq \{\alpha f + (1 - \alpha)L, \alpha g + (1 - \alpha)L\}$, it holds that

$$\alpha g + (1 - \alpha)L \in C(M) \Rightarrow \alpha f + (1 - \alpha)L \in C(M).$$

Let us denote by Γ^* the exogenous set of priors reflecting some objective ambiguity in the environment, and let us require that $\Gamma^* \subseteq \Delta(\Sigma)$ is compact and convex. Two more axioms are needed for the representation with exogenous priors.

S10. Γ^* -monotonicity:

$$\int_S \bar{u}(f(s)) dP(s) \geq \int_S \bar{u}(g(s)) dP(s) \quad \forall P \in \Gamma^* \Rightarrow f \triangleright_C g.$$

S11. Γ^* -ambiguity:

$$f \triangleright_C g \Rightarrow \int_S \bar{u}(f(s)) dP(s) \geq \int_S \bar{u}(g(s)) dP(s) \quad \forall P \in \Gamma^*.$$

The result about exogenous priors is then related to the one with endogenous priors through the axioms S10 and S11. Whenever the representation in terms of an endogenous set of priors Γ applies for a choice correspondence C , then

1. C satisfies Γ^* -monotonicity if and only if $\Gamma \subseteq \Gamma^*$.
2. C satisfies Γ^* -ambiguity if and only if $\Gamma^* \subseteq \Gamma$.

2.3 Ecological Rationality

Alternative to the account of rationality and decision making that we have seen so far, be it Bayesian or not, is the position that views DM as endowed with an *adaptive toolbox* of heuristics which are tailored for various and specific decision situations. This school of thought has the Nobel laureate Herbert Simon as prominent precursor (see e.g. [Simon, 1955], [Simon, 1990], [Simon, 1992]), and nowadays is centered around the research of Gerd Gigerenzer and colleagues at the Max Planck Institute for Human Development in Berlin. A cornerstone of this approach is that the rationality of a choice criterion has to be evaluated from its performance in specific decision problems, when competing with other possible criteria and heuristics. We largely sympathize with the idea of connecting the notion of rational choice to ecological and evolutionary considerations, and although our models will look considerably different than theirs from a formal point of view, we share similar intuitions about rationality from a general and conceptual point of view. It is then useful to examine this line of research in more detail.

The starting point of the explanation should probably specify that the representation of DM is no longer based on preferences, beliefs, attitudes, or similar internal attributes. Instead, DM possesses a toolbox of relevant heuristics to choose from when facing a decision problem. Heuristics are simple rules of thumb to make decisions. Unlike the results from the previous sections, a heuristic does not aim to be a universal solution for any optimization problem:

[...] heuristics do not try to optimize (i.e., find the best solution), but rather satisfice (i.e., find a good-enough solution). Calculating the maximum of a function is a form of optimizing; choosing the first option that exceeds an aspiration level is a form of satisficing. [Gigerenzer, 2008]

Heuristics are often deemed *fast and frugal* ([Gigerenzer and Goldstein, 1996]), because they may decide to ignore part of the information available. This relates to some important misconceptions that are strongly opposed by Gigerenzer and colleagues. Traditionally, a major justification for the use of heuristics has been based on the *accuracy-efficiency trade-off* argument: since the computations needed for optimizing are often intractable (NP-hard) in real-world problems, cognitive systems are allowed to rely on simpler heuristics (see for instance [Beach and Mitchell, 1978] and [Shah and Oppenheimer, 2011]). According to Gigerenzer's school, the implicit assumption behind this argument is that heuristics are always second-best strategies, that we use instead of logic or probability theory (which would be the first-best strategies) because of our cognitive limitations. Had we been provided with all the information, infinite time and unbounded computational power, there would be no reason for the existence of such heuristics. The sustainers of the adaptive toolbox firmly object to this view by

maintaining that there is no necessary trade-off between accuracy and efficiency, and contrast the accuracy-efficiency trade-off argument with the *less-is-more* argument. An example might help to understand this point and how heuristics work.

Suppose the problem at hand concerns how to invest money in N assets. Harry Markowitz proved that there exists an optimal portfolio that maximizes the returns and minimizes the risk ([Markowitz, 1952]), and later won the Nobel prize in economics for his work. A simple heuristic to solve the same allocation problem could be to invest equally in all the N assets. This is called the $1/N$ rule, and it is the strategy used by Markowitz himself for his retirement investments ([Gigerenzer and Brighton, 2009]). It may seem strange that Markowitz trusted this extremely simple heuristic more than the optimizing strategy that he himself designed and was awarded the Nobel prize for. However, his choice looks less bizarre in the light of some recent results on asset allocation in the financial market. [DeMiguel et al., 2009] compared fourteen optimal allocation policies with the $1/N$ rule in seven allocation problems, and showed that none of the fourteen optimizing strategies scored better than the simple heuristic on various financial measures. Although it can sound puzzling again, there are statistical reasons why this is the case: main causes for the success of the $1/N$ rule are the high unpredictability of the problem and the size of the learning sample (see [Gigerenzer and Brighton, 2009] for details). With respect to the latter, it must be stressed that the optimizing strategies could count on ten years of stock market data to estimate the parameters in the model, while the simple $1/N$ heuristic makes no use of past information at all. Still, the heuristic had superior performances. This is an instance of the *less-is-more* principle, that should cast doubt on the conception of heuristics as second-best options. In this case, disregarding relevant information and appealing to a computationally trivial strategy proved to be comparatively better than much more complex optimizing strategies. No tension between accuracy and efficiency showed up here.

For our purposes, the term “comparatively better” is of primary importance, and brings us back to the theme of ecological rationality and evolutionary success. Needless to say, the ecological and evolutionary success of a criterion essentially depends on the environment. Heuristics and choice rules that succeed in some environments or decision problems may be completely outperformed in others. What is at stake here is not the superiority of a simple heuristic over multiple regression models in general, but rather the question of which environments make simple heuristics more accurate than very sophisticated optimizing strategies. For example, the more unpredictable the environment and the smaller the learning sample, the bigger the advantage of the $1/N$ rule over the optimizing allocation strategies; but less so when the environment is more predictable and the learning sample size gets larger.

A central tenet of ecological rationality is that the notion of rational choice is defined by correspondence (how successful a choice is in a given environment)

rather than by coherence (whether a choice is consistent with respect to some axiomatic system). This account of rationality stands in opposition to the ones we have encountered in the previous sections: as [Gigerenzer, 2008] says,

Behavior is often called rational if and only if it follows the laws of logic or probability theory, and psychological research has consequently interpreted judgments that deviate from these laws as reasoning fallacies. From a Darwinian perspective, however, the goal of an organism is not to follow logic, but to pursue objectives in its environment, such as establishing alliances, finding a mate, and protecting offspring. Logic may or may not be of help. The rationality of the adaptive toolbox is not logical, but ecological; it is defined by correspondence rather than coherence.

Accordingly, a preponderant part of this program consists in studying the adaptive selection of heuristics, whose aim is to understand what environments make a certain heuristic ecologically rational. Once the scopes where the heuristic is pertinent have been uncovered, the analysis should focus on whether people actually adopt the heuristic where it is ecologically rational, and should refrain from testing if a given heuristic is used indistinctly. *Homo heuristicus* has not been endowed with a unique and universal principle fit for all practical purposes; *homo heuristicus* is an organism that reasons and acts conforming to multiple modular criteria, which he pulls out of his toolbox according to their performance in different environments.

To conclude, the paradigm of *homo heuristicus* has a threefold aspiration. A *descriptive* investigation, that shall describe the adaptive toolbox, its heuristics, as we observe them in use, and their building blocks. A *normative* claim, centered on the definition of ecological rationality, that shall determine the features of the environment favoring, and prescribing, a given heuristic. And, ultimately, the *design* of new heuristics and environments to improve decision making.

Chapter 3

Background on (Evolutionary) Game Theory

Trust everyone, but cut the cards. (F. P. Dunne)

This chapter aims at providing the tools which we will make use of in later results. While the preceding chapter was meant to be primarily conceptual, and introduced the ideas following the historical development of the field, here we will focus on the relevant game-theoretic notions and evolutionary models in a more compact and systematic way. We do not intend to present all the main achievements of game theory and evolutionary game theory, but rather to lay out the formal concepts that will be used in later investigations.

3.1 Game Theory

Game theory, as we know it nowadays, started with the seminal work of John von Neumann and Oskar Morgenstern ([von Neumann and Morgenstern, 1944]), where they also gave axiomatic foundation to expected utility maximization, as we have seen in Chapter 2. *Game theory* is the branch of mathematics that studies multi-agent interactive situations where the final consequences of DM's actions essentially depend on the actions chosen by other decision makers (called *players*). Examples are poker and chess, but also auctions and oligopolies. A more general and neutral way of naming game theory could thus be *interactive decision theory*. The crucial difference from decision theory is that when a player is aware of the interactive nature of the decision problem and realizes that the final outcome is determined by the choices of other players like her, she can try to put herself in the shoes of other intelligent agents and to anticipate their decisions to her advantage. This is the essence of *strategic thinking*.

The literature on game theory has grown at least as big as that in decision theory, but in this section we will focus only on *non-cooperative* game theory, and on *static* games in particular.¹

Definition 3.1. A *static game* is a tuple $G = \langle N, X, (A_i, u_i)_{i \in N}, \pi \rangle$, where

- $N = \{1, 2, \dots, n\}$ is the set of players;
- X is a set of *material* payoffs, or outcomes;
- A_i is the set of actions available to player i ;
- $\pi: \vec{A} \rightarrow X$ is the outcome function, with $\vec{A} := \prod_{i \in N} A_i$;
- $u_i: X \rightarrow \mathbb{R}$ is the *subjective* (vNM) utility function of player i .

In what follows, the sets N, X and A_i ($\forall i \in N$) are assumed to be finite.

Terminology and notation A few remarks, that will be important in the next chapters, are in order here. Following [Battigalli, 2016], the first four bullets of Definition 3.1 list what we call the *rules of the game*. The fifth bullet, the vNM utility function u_i , specifies the subjective preferences of player i over the possible outcomes. As we have seen in Chapter 2, the vNM utility function u_i reflects the agent's subjective valuation of the outcomes, and in many conceivable scenarios it is private information which is accessible to player i only.

We say that there is *complete information* in a game $G = \langle N, X, (A_i, u_i)_{i \in N}, \pi \rangle$ if it is common knowledge among all the players that G is the game that is being played. Complete information concerns the players' *interactive knowledge* about the rules of the game and their subjective preferences. Whenever the first four bullets of Definition 3.1 are commonly known, we say that there is common knowledge of the rules of the game. However, game G still features incomplete information unless the functions u_i are common knowledge too.

There is *perfect information* in a (dynamic) game when players move one at a time and all the players are informed of all the previous moves (chance moves included).² If perfect information fails, we say that the information in the game is *imperfect*. In particular, if a game has imperfect information because different players hold different pieces of information about past moves and/or the realization of chance moves, the game is said to have *asymmetric information*. It is of course possible that a game has imperfect and symmetric information, for

¹The reader can always consult [Battigalli, 2016], [Osborne and Rubinstein, 1994], or [Fudenberg and Tirole, 1991] for more details on all the game-theoretic notions presented in this section.

²Although featuring perfect or imperfect information is not a property of static games, it is useful to mention it here in order to stress the *qualitative* difference between perfect and complete information.

	<i>I</i>	<i>II</i>
<i>I</i>	1; 1	2; 5
<i>II</i>	5; 2	0; 0

Table 3.1: A static game.

instance if all the players are aware of all past moves except for the realization of a chance move. It must be clear that (im)perfect and (a)symmetric information are assumptions about the rules of the game, differently than (in)complete information, which is an assumption about the players' interactive knowledge. Presuming that players want to win, chess is an example of a game with complete and perfect information, whereas in poker there is complete but asymmetric information. On the other hand, auctions are examples of games with incomplete information if we assume that the participants' subjective valuations of items are not commonly known.

We call *profile* any list of objects (y_1, y_2, \dots, y_n) , where each y_i belongs to set Y_i , for $1 \leq i \leq n$. A typical profile is denoted by $\vec{y} := (y_1, \dots, y_n) \in \vec{Y} := \prod_{i \in N} Y_i$. For example, $\vec{a} := (a_1, \dots, a_n) \in \vec{A}$ is termed *action profile*, where each $a_i \in A_i$ is a possible action of player i . When we subtract the i^{th} element from a profile $\vec{y} = (y_1, \dots, y_i, \dots, y_n)$, we denote the reduced profile by $\vec{y}_{-i} := (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n) \in \vec{Y}_{-i} := \prod_{j \neq i} Y_j$. When we want to highlight the i^{th} component of the profile $\vec{y} = (y_1, \dots, y_i, \dots, y_n)$, we write the profile as (y_i, \vec{y}_{-i}) . Different subscripts denote that the elements come from different sets $(y_1, y_i, y_j, \dots) \in Y_1 \times Y_i \times Y_j \dots$, as opposed to different superscripts that mark the distinction between elements belonging to the same set $(y_i^1, y_i', y_i^i, y_i^j, \dots \in Y_i)$. For two-player games, we will often refer to player 1 as Ann and to player 2 as Bob.

Example 3.2. Consider the game defined by

- $N = \{1, 2\}$
- $A_1 = A_2 = \{I, II\}$
- $X = \vec{A} = A_1 \times A_2$
- $\pi(\vec{a}) = \vec{a}, \forall \vec{a} \in \vec{A}$
- $u_1(I, I) = u_2(I, I) = 1, u_1(I, II) = u_2(II, I) = 2, u_1(II, I) = u_2(I, II) = 5, u_1(II, II) = u_2(II, II) = 0.$

Table 3.1 depicts the game in its matrix form. Each entry in the (bi-)matrix is a pair of numbers, whose first component is the utility of player 1 (row player) and second component is the utility of player 2 (column player). This game is an example of an *anti-coordination* game, since each player prefers her action to differ from that of her co-player.

	<i>I</i>	<i>II</i>
<i>I</i>	1; 1	0; 0
<i>II</i>	0; 0	2; 2

Table 3.2: Hi-Lo game.

Example 3.3. Traveler's dilemma. A celebrated static game is the *Traveler's dilemma* ([Basu, 1994]). The story behind the game goes as follows. An airline company lost the suitcases of two customers, Ann (player 1) and Bob (player 2). As it happens, both suitcases were identical and contained equally valuable antiques. A manager of the airline company is assigned to the case and, being unable to figure out the precise value of the antiques, designs the following reimbursement policy. Ann and Bob are asked to separately claim an amount between \$2 and \$100. Since the antiques were equally valuable, if they both claim the same amount, then probably they are being honest and they will be reimbursed that amount. Otherwise, if one of them claims more than the other, the one who claimed the higher amount will get the lower amount minus \$2, and the one who claimed the lower amount will get what (s)he claimed plus \$2. Assuming that the subjective utility of Ann and Bob is linear in money, the traveler's dilemma is defined by:

- $N = \{1, 2\}$
- $A_1 = A_2 = \{\$2, \$3, \dots, \$100\}$
- $X = \{(x_1, x_2) \in \{\$2, \dots, \$100\} \times \{\$2, \dots, \$100\} : x_1 = x_2\} \cup \{(x_1, x_2) \in \{\$0, \dots, \$101\} \times \{\$0, \dots, \$101\} : x_2 = x_1 \pm \$4\}$
- $\pi(a_1, a_2) = \begin{cases} (a_1 + \$2, a_1 - \$2) & \text{if } a_1 < a_2 \\ (a_2 - \$2, a_2 + \$2) & \text{if } a_1 > a_2 \\ (a_1, a_2) & \text{if } a_1 = a_2 \end{cases}$
- $\forall (x_1, x_2) \in X, u_1(x_1, x_2) = x_1$ and $u_2(x_1, x_2) = x_2$.

Example 3.4. Hi-Lo. We can also represent a game by means of its matrix form directly, as in Table 3.2. If not specified otherwise, the numbers in the matrix express the subjective utilities of the players. The game in Table 3.2 is a *coordination* game, in that the players desire to coordinate on the same action, and it is often called *Hi-Lo* coordination game, because the coordination on one action (action *II*) gives all players a higher utility than the coordination on the other.

Definition 3.5. Symmetric game. Whenever a game is such that, for all players $i, j \in N$:

1. $A_i = A_j$, and
2. $u_i(\pi(a_i, \vec{a}_{-i})) = u_j(\pi(a_j, \vec{a}_{-j}))$, for $a_i = a_j$ and $\vec{a}_{-i} = \vec{a}_{-j}$,

we say that the game is *symmetric*.

In words, in a symmetric game all players have the same available actions and the utility attained by a player is independent of her role in the game. In fact, in symmetric games there are no distinctions by role. The player in role i has the same possible actions as the player in role j , and, for any fixed action profile of the co-players $\vec{a}_{-i} = \vec{a}_{-j}$, the utility to player i from a certain action is the same as that to player j from the same action. For these reasons, to define a symmetric game it suffices to specify only a generic action set $A = A_i$ for any player $i \in N$, and a generic utility function $u_i : \vec{A} \rightarrow \mathbb{R}$, where the set of all action profiles is given by $\vec{A} = A^{|N|}$. Consequently, we can also drop the index i from the function u_i by the convention that $u(\pi(a, a_2, \dots, a_n))$ denotes the utility of playing action $a \in A$ against the profile $(a_2, \dots, a_n) \in A^{|N|-1}$. The profile $(a_2, \dots, a_n) \in A^{|N|-1}$ will then be denoted by $\vec{a}_{-1} \in A^{|N|-1}$. A symmetric game G is thus specified by a tuple $G = \langle N, X, A, u, \pi \rangle$.

Note that all the games introduced in the examples of this section are symmetric. As it will be addressed later, symmetric games occupy an important position in evolutionary game theory, and will be the main objects of our analysis.

Solution concepts So far, we have seen how game theory enables us to formulate interactive decision problems in a precise mathematical language. But without further assumptions we would be unable to draw any conclusions about the behavior of the players. These assumptions pertain to the rationality and the beliefs of the players. Solution concepts implicitly carry these assumptions about the players' rationality and beliefs. We have talked at length about rationality and beliefs in Chapter 2, and the considerations from there also apply to the present context: players are nothing but (interactive) decision makers. In a general sense, to be made more precise, solution concepts represent fixpoints where rationality and beliefs of all players are mutually consistent. A bit more precisely, a solution concept gives a set of action profiles where the action of each player is consistent with his or her rationality and with his or her beliefs about the others' rationality, beliefs and actions. The most famous example is probably the Nash equilibrium.

Definition 3.6. Nash equilibrium. For a given game G , an action profile $\vec{a}^* \in \vec{A}$ is called a *Nash equilibrium* if

$$\forall i \in N, \forall a_i \in A_i, \quad u_i(\pi(a_i^*, \vec{a}_{-i}^*)) \geq u_i(\pi(a_i, \vec{a}_{-i}^*)).$$

Moreover, if the inequality holds strict for all players i , then \vec{a}^* is a *strict* Nash equilibrium. In symmetric games, a Nash equilibrium where all the players play the same action is said to be *symmetric*.

In plain words, an action profile \vec{a}^* is a Nash equilibrium if no player can profit by unilaterally deviating from her action in \vec{a}^* . The action choices in \vec{a}^* are thus mutually consistent and tied together.

Example 3.7. Example 3.2 continued. Consider the game from Example 3.2 again. In that game there are two Nash equilibria in pure actions: (I, II) and (II, I) . Since it is an anti-coordination game it should come as no surprise that players choose different actions when they are in equilibrium.

Example 3.8. Traveler's dilemma continued. We now want to find the profile consistent with Nash equilibrium in the Traveler's dilemma. To do that, consider a profile (a_1, a_2) where one of the two claims is bigger than the other, for instance $a_1 \geq a_2$. In this case, the action of Ann (player 1) is not a best reply to Bob's action: Ann's best reply would be to claim $a_2 - 1$. When Ann settles on her best reply, the resulting profile is (a'_1, a_2) , with $a'_1 = a_2 - 1$. But now Bob's action is not a best reply to Ann's action, and he should change his action to $a'_2 = a'_1 - 1 = a_2 - 2$. By continuing this reasoning all the way down, we realize that the only profile where each player is choosing a best reply to the other's action is $(\$2, \$2)$, which is consequently the only Nash equilibrium in the Traveler's dilemma. Ponder over this for a second: would you like to be consistent with Nash equilibrium here? That's why the title of the paper that first introduced the Traveler's dilemma is: *The Traveler's Dilemma: Paradoxes of Rationality in Game Theory* ([Basu, 1994]).

Example 3.9. Hi-Lo continued. In coordination games players want to coordinate on the same action, as reflected in the two equilibrium profiles of the Hi-Lo game: (I, I) and (II, II) . Both players would prefer to coordinate on the profile (II, II) , but note that both action profiles are consistent with Nash equilibrium.

Let us denote by \mathcal{G} the class of all static games. For a game $G \in \mathcal{G}$, let \vec{A}_G be the set of action profiles in G . Then we can formalize the notion of solution concept (for static games) by the following definition.

Definition 3.10. Solution concept. A *solution concept* is an element of the direct product $\prod_{G \in \mathcal{G}} 2^{\vec{A}_G}$. I.e., a solution concept is a function

$$F : \mathcal{G} \rightarrow \bigcup_{G \in \mathcal{G}} 2^{\vec{A}_G}$$

such that $\forall G \in \mathcal{G}, F(G) \subseteq \vec{A}_G$.

Mixed equilibria If any game had at least one Nash equilibrium in pure actions, we could imagine that after a process of mutual reasoning, mutual learning or similar, players are able to reach some stationary state, corresponding to a Nash equilibrium of the game. Unfortunately, there are games, like the one in the following table (known in the literature as the *rock-paper-scissors* game), that have no pure Nash equilibria.

	<i>I</i>	<i>II</i>	<i>III</i>
<i>I</i>	0; 0	-1; 1	1; -1
<i>II</i>	1; -1	0; 0	-1; 1
<i>III</i>	-1; 1	1; -1	0; 0

It may seem that players will be stuck in an infinite loop in such a game with no possibility of ever stabilizing on some equilibrium: if Ann chooses *I*, then Bob should choose *II*, but then Ann would prefer to switch to *III*, and so Bob would rather play *I*, to which Ann would respond *II*, that will make Bob play *III* and Ann reply with *I*, and so on. Fortunately, [Nash, 1950] found a way out of this predicament. Nash's solution involves the concept of mixed equilibrium, that is based on the *mixed extension* of a game.

Definition 3.11. Mixed extension. The *mixed extension* of a game $G = \langle N, X, (A_i, u_i)_{i \in N}, \pi \rangle$ is the game $\bar{G} = \langle N, \Delta(X), (\Delta(A_i), \bar{u}_i)_{i \in N}, \bar{\pi} \rangle$, where

- N is the set of players
- $\Delta(X) := \{L \in [0, 1]^X : \sum_{x \in X} L(x) = 1\}$ is the set of (lottery) outcomes
- $\Delta(A_i) := \{\alpha_i \in [0, 1]^{A_i} : \sum_{a_i \in A_i} \alpha_i(a_i) = 1\}$ is i 's set of (mixed) actions
- $\bar{\pi} : \prod_{j \in N} \Delta(A_j) \rightarrow \Delta(X)$ is the expected outcome function, that associates with each $\vec{\alpha} \in \prod_{j \in N} \Delta(A_j)$ the lottery $\bar{\pi}(\vec{\alpha}) = L$ such that, for $\pi(\vec{\alpha}) \in X$,

$$L(\pi(\vec{\alpha})) = \prod_{j \in N} \alpha_j(a_j).$$

- $\bar{u}_i : \Delta(X) \rightarrow \mathbb{R}$ is i 's subjective expected utility function, such that

$$\bar{u}_i(\bar{\pi}(\vec{\alpha})) := \sum_{\vec{a} \in \vec{A}} u_i(\pi(\vec{a})) \prod_{j \in N} \alpha_j(a_j).$$

The definition of $\bar{\pi}$ highlights that the mixed actions of the players are implicitly assumed to be *statistically independent*. Moreover, the mixed extension of a symmetric game is a symmetric game itself. Building on Definition 3.11, we are now in the position of defining mixed Nash equilibria.

Definition 3.12. Mixed Nash equilibrium. A profile of mixed actions $\vec{\alpha}^* \in \prod_{i \in N} \Delta(A_i)$ is a *mixed Nash equilibrium* of a game G if it is a Nash equilibrium of the mixed extension \bar{G} of G .

Nash proved that any finite game (i.e., a game where the sets N and A_i , for all $i \in N$, are finite) admits at least one mixed Nash equilibrium. Obviously, pure actions can be viewed as degenerate mixed actions, so that $A_i \subset \Delta(A_i)$, $\forall i \in N$. Therefore, any Nash equilibrium in pure actions \vec{a}^* corresponds to a mixed Nash equilibrium $\vec{\alpha}^*$, where $\alpha_i^*(a_i^*) = 1$, and $\alpha_i^*(a_i) = 0$ for $a_i \neq a_i^*$ and $i \in N$.

Since the rock-paper-scissors game is a finite game, it has a Nash equilibrium in mixed actions, when both Ann and Bob play each action $\frac{1}{3}$ of the times. I.e., the unique mixed Nash equilibrium of the rock-paper-scissors is the mixed action profile (α_1^*, α_2^*) , with $\alpha_i^*(I) = \alpha_i^*(II) = \alpha_i^*(III) = \frac{1}{3}$ for $i = 1, 2$.

Example 3.13. Example 3.2 continued. In the game of Example 3.2 there are three mixed Nash equilibria. Apart from the two pure equilibria already identified in Example 3.7, there is a genuine mixed Nash equilibrium \vec{a}^* such that $\alpha_i^*(I) = \frac{1}{3}$ and $\alpha_i^*(II) = \frac{2}{3}$, for $i = 1, 2$. Notice that the expected utility of the players at the mixed equilibrium is $\frac{5}{3}$, which is strictly less than what either of them would get if they played any of the other two equilibria in pure actions.

Example 3.14. Hi-Lo continued. Consider again the Hi-Lo game of Table 3.2. Apart from the two Nash equilibria in pure actions, there is another Nash equilibrium in mixed actions, where both players choose I with probability $\frac{2}{3}$ and II with probability $\frac{1}{3}$. For both players, the expected utility at the mixed equilibrium is $\frac{2}{3}$, which is again strictly less than the expected utility at any of the pure equilibria.

Finally, we want to say a few words about the interpretation of mixed actions. To fix ideas, let us suppose that we are dealing with a two-player game. In classic game theory, the mixed action α_1 is interpreted as Ann randomizing over her own pure actions. Yet, from experimental psychology we know that humans are very poor and imprecise randomization devices, and seem incapable of implementing the play recommended by a given mixed action. Real players constantly deviate from the given frequencies. In accordance with this observation, *epistemic* game theory assumes that players can only choose pure actions: players never randomize. The reader might now wonder if we are falling back into the problem of having games without any equilibrium. After all, we were brought to considering mixed actions in order to solve that impasse. The good news is that the epistemic program in game theory can maintain both the assumption about pure actions and the existence of equilibria in any finite game. The solution lies in the interpretation of mixed equilibria as *equilibria in beliefs*. Concretely, the mixed action α_1 of Ann is no longer a randomized action of Ann (players never randomize), but rather a belief of Bob about what Ann is going to play. In the rock-paper-scissors game, for example, both Ann and Bob can only choose from the set $\{I, II, III\}$, and their beliefs are in equilibrium when each player believes any of the co-player's actions with probability $\frac{1}{3}$, and believes that the co-player believes the same. Indeed, when Ann believes any of Bob's actions to be equiprobable, any action is

a best reply for her; and similarly for Bob. So, if Bob believes that Ann is going to play *I*, *II* or *III* with equal probability, and that she also believes that he is going to play *I*, *II* or *III* with equal probability, then Bob believes that Ann is rational. And, mutatis mutandis, Ann believes that Bob is rational. Therefore, in such a situation both players believe in rationality and believe that the co-player expects any action with equal probability. These epistemic conditions give rise to an equilibrium in beliefs that corresponds to the mixed Nash equilibrium of the game (see [Brandenburger, 1992], [Aumann and Brandenburger, 1995]). The solution concept of (mixed) Nash equilibrium is thus characterized and derived from epistemic assumptions about players' (interactive) beliefs and rationality. Under this interpretation players only choose pure actions, and the mixture over one's own actions is projected onto the probabilistic belief of the co-player.

There are at least two main formal structures in the literature to reason about the interactive rationality of the players: type spaces and (multi-agent) Kripke models. A logical analysis of these structures will be the topic of Chapter 4. It has to be stressed, however, that the notion of (epistemic) types that we will consider in Chapter 4 is different than the notion of types that will be used in the rest of this thesis. Epistemic types are essentially hierarchies of interactive beliefs (see [SectionChapterEmiliano with the definition]). The types that we will introduce in later chapters, instead, are the types of the players existing in a population, as customary in evolutionary game theory. Since our focus will be the evolution of different ways of making choices, these types will consist of a subjective utility function, a subjective belief, and a decision criterion.

In general, epistemic game theory holds that the strategic analysis of a game should start from the players' beliefs. For a game $G = \langle N, X, (A_i, u_i), \pi \rangle$, let $P_i \in \Delta(A_{-i})$ denote the (Bayesian) belief of player i about the actions of her co-players, and let $\text{supp}\alpha_i \subseteq A_i$ denote the support of mixed action α_i . The following lemma justifies the soundness of such an epistemic approach to game theory, as well as the claim that players never randomize.³

Lemma 3.15. *For a game $G = \langle N, X, (A_i, u_i)_{i \in N}, \pi \rangle$ and a belief $P_i \in \Delta(A_{-i})$ of player i , the following are equivalent:*

1. $\alpha_i^* \in \text{argmax}_{\alpha_i \in \Delta(A_i)} \sum_{\vec{a} \in \vec{A}} u_i(\pi(\vec{a})) \alpha_i(a_i) P_i(\vec{a}_{-i})$,
2. $\text{supp}\alpha_i^* \subseteq \text{argmax}_{\alpha_i \in \Delta(A_i)} \sum_{\vec{a} \in \vec{A}} u_i(\pi(\vec{a})) \alpha_i(a_i) P_i(\vec{a}_{-i})$.

In words, Lemma 3.15 says that the players, whatever their beliefs about the co-players are, cannot increase their expected utility by playing a mixed action. For any possible belief P_i , there is always a pure action $a_i \in A_i$ that maximizes expected utility: players never randomize, and they never need to.

³The proof is not difficult, but we are not going to present it here for reasons of succinctness. The interested reader can find it in [Battigalli, 2016].

In what follows, we assume that players can always choose pure actions only, in line with the approach supported by epistemic game theory.

A solution concept that will play a major role in later chapters is the *Bayesian Nash equilibrium*. Before presenting it, we have to explain what Bayesian games are.

Bayesian games and Bayesian Nash equilibrium The key to understanding Bayesian games is incomplete information. We said that complete information is an assumption about the players' interactive knowledge of the game, and whenever the game is not commonly known, players have incomplete information of the game. For example, players may lack (common) knowledge about the others' preferences, as well as about the outcome function. A game that features incomplete information gives rise to a Bayesian game, in that players need to form (Bayesian) beliefs about all the relevant unknown variables. *Bayesian games* are the mathematical structures that permit to formalize games where the information is not complete. The main ideas to deal with games featuring this augmented uncertainty come from the seminal work by John Harsanyi ([Harsanyi, 1967], [Harsanyi, 1968a], and [Harsanyi, 1968b]), who enriched game structures with a set of possible worlds and possible types for each player.

Definition 3.16. Bayesian game. A Bayesian game is a structure $BG = \langle N, S, X, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, P_i, u_i)_{i \in N}, \pi \rangle$, where

- N is the set of players
- S is a set of possible states of the world, assumed finite for simplicity
- X is the set of material payoffs, or outcomes
- Θ_i is the set of (subjective) *utility types* of player i
- T_i is the set of *types* of player i
- A_i is the set of actions available to player i
- $\tau_i : S \rightarrow T_i$ is the signal function, or *type function*, of player i
- $\vartheta_i : T_i \rightarrow \Theta_i$ is the *utility type function* of player i
- $P_i \in \Delta(S)$ is player i 's prior belief, such that $\forall t_i \in T_i, P_i(\tau_i^{-1}(t_i)) > 0$
- $u_i : \Theta_i \times X \rightarrow \mathbb{R}$ is player i 's subjective utility function
- $\pi : \vec{\Theta} \times \vec{A} \rightarrow X$ is the outcome function, with $\Theta := \prod_{i \in N} \Theta_i$.

$P_i(s_1) = \frac{1}{8}$	t_2^1	t_2^2	$P_i(s_2) = \frac{1}{4}$												
	I	I													
	II	II													
t_1^1	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="text-align: center; padding: 2px;">I</td><td style="padding: 2px;">1;1</td><td style="padding: 2px;">2;5</td></tr> <tr><td style="text-align: center; padding: 2px;">II</td><td style="padding: 2px;">5;2</td><td style="padding: 2px;">0;0</td></tr> </table>	I	1;1	2;5	II	5;2	0;0	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="text-align: center; padding: 2px;">I</td><td style="padding: 2px;">1;1</td><td style="padding: 2px;">2;0</td></tr> <tr><td style="text-align: center; padding: 2px;">II</td><td style="padding: 2px;">5;0</td><td style="padding: 2px;">0;2</td></tr> </table>	I	1;1	2;0	II	5;0	0;2	
I	1;1	2;5													
II	5;2	0;0													
I	1;1	2;0													
II	5;0	0;2													
	I	I													
	II	II													
t_1^2	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="text-align: center; padding: 2px;">I</td><td style="padding: 2px;">1;1</td><td style="padding: 2px;">0;5</td></tr> <tr><td style="text-align: center; padding: 2px;">II</td><td style="padding: 2px;">0;2</td><td style="padding: 2px;">2;0</td></tr> </table>	I	1;1	0;5	II	0;2	2;0	<table style="width: 100%; border-collapse: collapse;"> <tr><td style="text-align: center; padding: 2px;">I</td><td style="padding: 2px;">1;1</td><td style="padding: 2px;">0;0</td></tr> <tr><td style="text-align: center; padding: 2px;">II</td><td style="padding: 2px;">0;0</td><td style="padding: 2px;">2;2</td></tr> </table>	I	1;1	0;0	II	0;0	2;2	
I	1;1	0;5													
II	0;2	2;0													
I	1;1	0;0													
II	0;0	2;2													
$P_i(s_3) = \frac{1}{4}$			$P_i(s_4) = \frac{3}{8}$												

Figure 3.1: A Bayesian game.

A few words and examples will help to clarify the meaning of the items in Definition 3.16. As noticed before, the most likely cause of incomplete information is the failure of common knowledge about the players' subjective utilities. The uncertainty about the others' subjective preferences can be expressed by means of the utility functions u_i that are now parameterized via the utility types $\theta_i \in \Theta_i$. The following example shows this point.

Example 3.17. Consider the Bayesian game depicted in Figure 3.1, defined by the following bullets:

- $N = \{1, 2\}$
- $S = \{s_1, s_2, s_3, s_4\}$
- $X = \vec{A}$
- $\Theta_i = T_i = \{t_i^1, t_i^2\} \quad \forall i \in N$
- $A_i = \{I, II\} \quad \forall i \in N$
- $\tau_1(s_1) = \tau_1(s_2) = t_1^1, \tau_1(s_3) = \tau_1(s_4) = t_1^2$
- $\tau_2(s_1) = \tau_2(s_3) = t_2^1, \tau_2(s_2) = \tau_2(s_4) = t_2^2$
- $\vartheta_i(t_i) = t_i \quad \forall i \in N, \forall t_i \in T_i$

- $P_i(s_1) = \frac{1}{8}, P_i(s_2) = P_i(s_3) = \frac{1}{4}, P_i(s_4) = \frac{3}{8} \quad \forall i \in N$
- $u_i(t_i^1, I, I) = 1, u_i(t_i^1, I, II) = 2, u_i(t_i^1, II, I) = 5, u_i(t_i^1, II, II) = 0,$
 $u_i(t_i^2, I, I) = 1, u_i(t_i^2, I, II) = 0, u_i(t_i^2, II, I) = 0, u_i(t_i^2, II, II) = 2, \forall i \in N$
- $\pi(\vec{\theta}, \vec{a}) = \vec{a} \quad \forall \vec{\theta} \in \vec{\Theta}, \forall \vec{a} \in \vec{A}.$

Intuitively, Ann is uncertain whether Bob's subjective utilities correspond to those of the game in Example 3.2 or to those of the Hi-Lo game in Example 3.4. Bob's uncertainty about Ann's utilities is exactly the same. Formally, this uncertainty is expressed by assigning two (utility-)types to each player. Type t_i^1 's utilities are those of Table 3.1, while type t_i^2 's utilities are those of Table 3.2.

Another familiar example of a game without common knowledge of the subjective utility functions is the popular board game (*Secret Mission*) *Risk*, where the players have to draw a mission card each from a deck of possible missions, without revealing to the others which mission they have been assigned to. Lack of knowledge about the others' subjective preferences is not the unique possible source of incomplete information in games, albeit presumably the most common. The next example, borrowed from [Battigalli, 2016], considers two agents involved in the production of a public good.

Example 3.18. Ann and Bob have to decide how much effort to invest in the production of a public good. Each of them can choose to put a quantity of effort $a_i \in [0, 1]$. The output y depends on their efforts according to the Cobb-Douglas production function:

$$y = K(a_1)^{\theta_1}(a_2)^{\theta_2},$$

where K is a constant parameter. The cost of the effort for each player is $c_i(a_i) = (a_i)^2$. The outcome function $\pi : [0, 1]^2 \rightarrow \mathbb{R}_+^3$, that expresses both the output and the costs of the production, is given by

$$\pi(\theta_1, \theta_2, a_1, a_2) = (K(a_1)^{\theta_1}(a_2)^{\theta_2}, (a_1)^2, (a_2)^2).$$

The subjective utility for each player i and each θ_i is given by

$$u_i(\theta_i, \pi(\theta_1, \theta_2, a_1, a_2)) = K(a_1)^{\theta_1}(a_2)^{\theta_2} - (a_i)^2 = y - c_i.$$

If the players do not know the parameters K, θ_1, θ_2 , then they do not know the material payoff function π , even if they would know the utility functions: $u_i(\theta_i, y, c_1, c_2) = y - c_i$, for all i and all θ_i . Indeed, even if Ann could know the parameter K and her own productivity θ_1 , she might not know how productive Bob is when he decides to put effort a_2 ; and similarly for Bob. Since the outcome depends on the productivity of both, the outcome function π is not known and the game features incomplete information.

Terminology Whenever there is common knowledge of the material payoff function (i.e., $\pi(\vec{\theta}, \vec{a}) = \pi(\vec{\theta}', \vec{a})$ for all $\vec{\theta}, \vec{\theta}' \in \vec{\Theta}$), and for all players i , the utility function u_i depends only on the action profile \vec{a} and on i 's own utility type θ_i , the game is said to have *private values*. This is the case of the game in Example 3.17. If the function π is not common knowledge (i.e., $\pi(\vec{\theta}, \vec{a}) \neq \pi(\vec{\theta}', \vec{a})$ for some $\vec{\theta}, \vec{\theta}' \in \vec{\Theta}$), the game is said to have *interdependent values*. The public good production game of Example 3.18 has interdependent values. Usually, it is also assumed that all the players hold the same prior probability over the states of the world, $\forall i \in N \ P_i = P$. In this case, the game is said to have *common prior*. Finally, the structure of the Bayesian game (i.e., all the items in Definition 3.16) is implicitly assumed to be commonly known by the players.

As already mentioned when we introduced Nash equilibrium, there is nothing in the definition of a game that describes or prescribes any specific behavior of the players. The same holds for Bayesian games as well: we need to make assumptions regarding the rationality and beliefs of the players in order to draw conclusions about their behavior. The concept of Bayesian Nash equilibrium is the equivalent of Nash equilibrium for Bayesian games.

Definition 3.19. Bayesian Nash equilibrium. A *Bayesian Nash equilibrium* of a Bayesian game BG is a profile of *policy functions* $(\sigma_i^* : T_i \rightarrow A_i)_{i \in N}$ such that, $\forall i \in N, \forall t_i \in T_i$,

$$\sigma_i^*(t_i) \in \operatorname{argmax}_{a_i \in A_i} \sum_{s \in S} P_i(s|t_i) u_i(\vartheta_i(t_i), \pi(\vartheta_i(t_i), \vec{\vartheta}_{-i}(\vec{\tau}_{-i}(s)), a_i, \vec{\sigma}_{-i}(\vec{\tau}_{-i}(s))))),$$

where $\vec{\tau}_{-i}(s) = (\tau_j(s))_{j \neq i} := \prod_{j \neq i} \{\tau_j(s)\}$, and similarly $\vec{\vartheta}_{-i}(\vec{t}_{-i}) := (\vartheta_j(t_j))_{j \neq i}$ and $\vec{\sigma}_{-i}(\vec{t}_{-i}) := (\sigma_j(t_j))_{j \neq i}$.

To illustrate the connection between Nash equilibrium and Bayesian Nash equilibrium, we have to introduce two complete information games associated with a given Bayesian game: the *ex ante* and the *interim* strategic form of the game.

Definition 3.20. Ex ante strategic form. Given a Bayesian game BG , we define the *ex ante* strategic form of BG the complete information game $G^{ex} = \langle N, X^{\vec{T}}, (\Sigma_i, u_i^{ex})_{i \in N}, \pi^{ex} \rangle$, where N and X are as in BG , and

- $\Sigma_i := A_i^{T_i}$ (i.e., the set of all functions from T_i to A_i) is i 's action set
- π^{ex} is the ex ante material payoff (or outcome) function, and u_i^{ex} is the ex ante utility function of player i , such that the function $u_i^{ex}(\pi^{ex}(\cdot)) : \vec{\Sigma} \rightarrow \mathbb{R}$ is defined by:

$$u_i^{ex}(\pi^{ex}(\vec{\sigma})) := \sum_{s \in S} P_i(s) u_i(\vartheta_i(\tau_i(s)), \pi(\vec{\vartheta}(\vec{\tau}(s)), \vec{\sigma}(\vec{\tau}(s))))$$

$$= \sum_{t_i \in T_i} P_i(t_i) \sum_{\vec{t}_{-i} \in \vec{T}_{-i}} P_i(\vec{t}_{-i} | t_i) u_i(\vartheta_i(t_i), \pi(\vartheta_i(t_i), \vec{\vartheta}_{-i}(\vec{t}_{-i}), \sigma_i(t_i), \vec{\sigma}_{-i}(\vec{t}_{-i}))).$$

In the ex ante strategic form all the uncertainty is considered from an ex ante perspective, when the players have not yet received any signal about their type. Each player is then asked to choose a policy function from Σ_i that specifies an action choice for any possible type of his or hers. The utility u_i^{ex} of player i in the ex ante version of the game corresponds to the expected utility given by the function u_i with respect to i 's probability P_i over the states, together with the policy function σ_i chosen by player i and the policy functions $\vec{\sigma}_{-i}$ chosen by the other players. The next remark then follows.

Remark 3.21. A profile $(\sigma_i^*)_{i \in N}$ is a Bayesian Nash equilibrium of the Bayesian game BG if and only if it is a Nash equilibrium of the ex ante strategic form G^{ex} of BG .

Definition 3.22. Interim strategic form. Given a Bayesian game BG , we define the *interim* strategic form of BG the complete information game $G^{in} = \langle \bigcup_{i \in N} T_i, X^{\vec{T}}, (A_{t_i}, u_{t_i})_{i \in N, t_i \in T_i}, \pi^{in} \rangle$, where

- $\bigcup_{i \in N} T_i$ is the set of players
- X is as in BG
- $A_{t_i} = A_i, \forall i \in N, \forall t_i \in T_i$ is the action set of player t_i
- π^{in} is the interim outcome function, and u_{t_i} is the interim utility function of player t_i , such that the function $u_{t_i}(\pi^{in}(\cdot)) : \prod_{j \in N} \prod_{t_j \in T_j} A_{t_j} \rightarrow \mathbb{R}$ is defined by:

$$u_{t_i}(\pi^{in}((a_{t_j})_{j \in N, t_j \in T_j})) := \sum_{\vec{t}_{-i} \in \vec{T}_{-i}} P_i(\vec{t}_{-i} | t_i) u_i(\vartheta_i(t_i), \pi(\vartheta_i(t_i), \vec{\vartheta}_{-i}(\vec{t}_{-i}), a_{t_i}, (a_{t_j})_{j \neq i})).$$

In the interim strategic form, the intuition is that the players have already received the signal about their type, and each type t_i is now an independent player that chooses its own action a_{t_i} and has its own utility function u_{t_i} . Note that the utility u_{t_i} in the interim version corresponds to the expected utility given by the utility function u_i with respect to the conditional probability $P_i(\cdot | t_i)$, together with the actions $a_{t_i} \in A_i$ and $(a_{t_j})_{j \neq i} \in \vec{A}_{-i}$. As before, we remark the following point.

Remark 3.23. A profile $(\sigma_i^*)_{i \in N}$ is a Bayesian Nash equilibrium of the Bayesian game BG if and only if the profile $(a_{t_i}^*)_{i \in N, t_i \in T_i}$, such that $a_{t_i}^* = \sigma_i^*(t_i)$ for all $i \in N$ and $t_i \in T_i$, is a Nash equilibrium of the interim strategic form G^{in} of BG .

Example 3.24. Example 3.17 continued. The following table represents the players' utilities of the ex ante strategic form of the Bayesian game from Example 3.17. The first component of ordered pairs like (I, I) specifies the action of type t_i^1 , and the second component specifies the action of type t_i^2 .

	I, I	I, II	II, I	II, II
I, I	1; 1	$7/8; 3/8$	$7/8; 5/2$	$3/4; 15/8$
I, II	$3/8; 7/8$	$11/8; 11/8$	$11/8; 7/8$	2; $11/8$
II, I	$5/2; 7/8$	$7/8; 11/8$	$13/8; 13/8$	0; $7/4$
II, II	$15/8; 3/4$	$11/8; 2$	$7/4; 0$	$5/4; 5/4$

There are three Nash equilibria in this game: $((I, II), (I, II))$, $((I, II), (II, II))$, and $((II, II), (I, II))$. The first corresponds to the profile of policy functions (σ_1, σ_2) such that $\sigma_1(t_1^1) = I, \sigma_1(t_1^2) = II, \sigma_2(t_2^1) = I, \sigma_2(t_2^2) = II$. Similarly, the second corresponds to the profile (σ_1, σ_2) such that $\sigma_1(t_1^1) = I, \sigma_1(t_1^2) = II, \sigma_2(t_2^1) = II, \sigma_2(t_2^2) = II$; the third corresponds to the profile (σ_1, σ_2) such that $\sigma_1(t_1^1) = II, \sigma_1(t_1^2) = II, \sigma_2(t_2^1) = I, \sigma_2(t_2^2) = II$.

The utilities for the interim strategic form are given by the following formulas. Since the utility $u_{t_i^1}$ of player $t_i^1 \in T_i$ is not dependent on the actions of the other type $t_i^2 \in T_i$ (and vice versa), we can ease notation by using formulas like $u_{t_i^1}(\pi^{in}(I, (I, II)))$ as an abbreviation for $u_{t_i^1}(\pi^{in}((I, II), (I, II)))$, and $u_{t_i^2}(\pi^{in}(II, (I, II)))$ for $u_{t_i^2}(\pi^{in}((I, II), (I, II)))$, where the first action is the choice of the player whose utility we are considering and the following pair specifies the actions of the co-players t_{3-i}^1 and t_{3-i}^2 . The interim utilities of player t_1^1 are:

$$\begin{aligned}
u_{t_1^1}(\pi^{in}(I, (I, I))) &= 1 & u_{t_1^1}(\pi^{in}(II, (I, I))) &= 5 \\
u_{t_1^1}(\pi^{in}(I, (I, II))) &= 5/3 & u_{t_1^1}(\pi^{in}(II, (I, II))) &= 5/3 \\
u_{t_1^1}(\pi^{in}(I, (II, I))) &= 4/3 & u_{t_1^1}(\pi^{in}(II, (II, I))) &= 10/3 \\
u_{t_1^1}(\pi^{in}(I, (II, II))) &= 2 & u_{t_1^1}(\pi^{in}(II, (II, II))) &= 0
\end{aligned}$$

Analogously, the interim utilities of player t_1^2 are:

$$\begin{aligned}
u_{t_1^2}(\pi^{in}(I, (I, I))) &= 1 & u_{t_1^2}(\pi^{in}(II, (I, I))) &= 0 \\
u_{t_1^2}(\pi^{in}(I, (I, II))) &= 2/5 & u_{t_1^2}(\pi^{in}(II, (I, II))) &= 6/5 \\
u_{t_1^2}(\pi^{in}(I, (II, I))) &= 3/5 & u_{t_1^2}(\pi^{in}(II, (II, I))) &= 4/5 \\
u_{t_1^2}(\pi^{in}(I, (II, II))) &= 0 & u_{t_1^2}(\pi^{in}(II, (II, II))) &= 2
\end{aligned}$$

The interim utilities of players t_2^1 are the same as those of player t_1^1 , and the interim utilities of t_2^2 are the same as those of t_1^2 . We can notice that when the co-players are playing actions (I, II) , both I and II are best replies for player t_1^1 and II is the unique best reply for player t_1^2 . Since the utilities of t_2^1 and t_2^2 equal those of t_1^1 and t_1^2 , the same holds for players t_2^1 and t_2^2 . Consequently, $((I, II), (II, II))$, $((II, II), (I, II))$, and $((I, II), (I, II))$ are Nash equilibria of the interim form of the game. Notice moreover that there is no other Nash equilibrium in the game.

For instance, if we consider the profile $((I, I)(II, I))$, we can observe that player t_1^1 would switch to action II , and so would player t_1^2 , for a resulting profile of $((II, II)(II, I))$. But then player t_2^1 would prefer to play action I rather than II , and player t_2^2 would prefer to switch from I to II , giving rise to the profile $((II, II)(I, II))$, where the players are finally in equilibrium. The three Nash equilibria that we have found in the interim strategic form correspond to the three policy functions obtained from the ex ante analysis above, as anticipated in Remark 3.23.

In later chapters, incomplete information games will play a crucial role. We already argued that common knowledge of subjective preferences is not a realistic assumption in many cases. It is even less realistic when we deal with evolutionary models and large populations of players, as customary in evolutionary game theory. There, a player is normally supposed to recurrently play a fixed game with different co-players randomly selected from the population at any repetition of the game. If we think of real-life scenarios, it might be sometimes possible that we know the subjective preferences of the co-players that we are matched with. But if we had to meet members of a population at random, most likely we would have no idea about our co-players' preferences in the majority of the cases. In all those circumstances, the situation we would be facing corresponds to a game with incomplete information.

3.2 Evolutionary Game Theory

The combination of game theory with evolutionary models and dynamical systems has proven to be very fruitful, and offered researchers many new insights over the last decades. As a result, game theory, in its population version, is nowadays a well-established instrument for analysis in many fields: biology, ecology, sociology, linguistics, ethology, philosophy, anthropology and others (e.g., [Skyrms, 1996], [Sinervo and Lively, 1996], [Bergstrom and Godfrey-Smith, 1998], [Skyrms, 2010], and [Franke, 2012]). The success gained by evolutionary game theory in all these areas stems from the possibility of studying (the attainability and the stability of) biological as well as sociological and linguistic traits as the outcome of the evolutionary selection between different and competing traits.

In evolutionary game theory, players are traditionally thought of and modeled as very primitive and unsophisticated organisms, like mindless animals or plants, whose behavior is mechanically determined by a genetic code that they received from their parents and will transmit to their offspring. The evolutionary success is defined in terms of differential reproduction: the more successful a behavior, the larger the offspring in the next generation. A behavior, trait or action under consideration and under selection is generically termed *phenotype*.

There are two different approaches to an evolutionary investigation of the interactions between members of a population. The *static* analysis is mainly

interested in the stability of a given state, trait or behavior. For this reason, it is also called *equilibrium* analysis. On the other side, the *dynamic* analysis focuses more on the attainability of a given state, trait or behavior. We will present both these methodologies, but we refer to [Huttegger and Zollman, 2013] for a more extensive discussion.

Static analysis The point of departure for the static analysis of population interactions is the concept of *evolutionarily stable strategy (ESS)* (see the work by [Maynard Smith and Price, 1973], [Maynard Smith, 1974], [Maynard Smith, 1982]).⁴ In order to introduce ESSs, we first need to define the unit of measure for evolutionary success. This quantity is called evolutionary payoff, or objective (as opposed to subjective) utility, or simply *fitness*. The fitness Φ of a behavior is specified by fixing the *fitness game*.

Definition 3.25. Fitness game. A *fitness game* $G = \langle N, X, (A_i, \Phi_i)_{i \in N}, \pi \rangle$ is a game where N, X, A_i and π are as in Definition 3.1, and $\Phi_i : X \rightarrow \mathbb{R}$ is the *fitness* function of player i .

Note that the functions Φ_i have the same domain and codomain as the functions u_i of Definition 3.1, but the interpretation is different. While u_i represents player i 's subjective utility (or happiness), which is an intimate feeling, the function Φ_i gives the fitness of player i , which is supposed to be an objective quantity.⁵

The notion of ESS is classically defined for two-player games and, strictly speaking, it only applies to symmetric games. However, it is customary in evolutionary game theory to talk about ESSs in asymmetric games too. In the latter case, the concept of ESS is referred to the symmetrized version of the game. This will not be an important issue here, since we will deal with symmetric two-player games for most of the results. A fitness game is *symmetric* if it satisfies Definition 3.5, where, for all $i \in N$, the function u_i is replaced by the function Φ_i .

Definition 3.26. ESS. Let $G = \langle N, X, A, \Phi, \pi \rangle$ be a symmetric two-player fitness game. A strategy (i.e., phenotype) $a^* \in A$ is *evolutionarily stable* if, for all $a \neq a^*$, it holds that:

1. $\Phi(\pi(a^*, a^*)) \geq \Phi(\pi(a, a^*))$;
2. $\Phi(\pi(a^*, a^*)) = \Phi(\pi(a, a^*)) \Rightarrow \Phi(\pi(a^*, a)) > \Phi(\pi(a, a))$.

⁴The terms “action”, “act”, and “strategy” are all synonymous in static games. The reader must also keep in mind that static games (as defined in Section 3.1) and the static analysis of evolutionary game theory are completely different notions.

⁵If we consider market competition for instance, an economist would maybe think of this quantity as money, with the intuition that only firms that maximize profits will survive in the long run, independent of the subjective happiness that each firm attaches to (different amounts of) money.

Whenever the consequent in point 2 is just a weak inequality, a^* is said to be a *neutrally stable strategy (NSS)*.

Intuitively, ESS captures the evolutionary stability of a monomorphic population. A population is termed *monomorphic* when all its players display the same phenotype. We call a -monomorphic a monomorphic population whose unique phenotype is $a \in A$. If the unique phenotype of a monomorphic population is an ESS of the fitness game, then there are no other phenotypes that can invade and thrive in the population. Indeed, imagine an a^* -monomorphic population confronting a few invaders playing action $a \in A$. Since the invaders represent only a minimal share of the population, each player will encounter the incumbent phenotype a^* most of the times. If condition 1 of Definition 3.26 holds strictly, there is no chance for the invaders to achieve a higher (expected) fitness than the incumbent type. Given that the number of offspring is determined by the fitness of a phenotype, a will be weeded out of the population in the next generations. If instead $\Phi(\pi(a^*, a^*)) = \Phi(\pi(a, a^*))$, the tie-breaking condition 2 of Definition 3.26 is crucial, since it compares the performance of each phenotype when paired with one of the few invaders. If the incumbent type a^* has higher fitness against the invader a than a against itself, then a^* has higher expected fitness in the population in general, and will be evolutionarily stable.

Example 3.27. Hi-Lo continued. Suppose a II -monomorphic population of players is recurrently playing the Hi-Lo fitness game, i.e., the fitness game $G = \langle N, X, A, \Phi, \pi \rangle$ such that

- $N = \{1, 2\}$
- $X = \vec{A}$
- $A = \{I, II\}$
- $\Phi(I, I) = 1, \Phi(I, II) = 0 = \Phi(II, I), \Phi(II, II) = 2$
- $\pi(\vec{a}) = \vec{a}, \forall \vec{a} \in \vec{A}$.

Since it holds that $\Phi(\pi(II, II)) = 2 > 0 = \Phi(\pi(I, II))$, action II is an ESS by condition 1 of Definition 3.26. Analogously, it also holds that $\Phi(\pi(I, I)) = 1 > 0 = \Phi(\pi(II, I))$, so that a I -monomorphic population cannot be invaded by phenotype II . Hence, both strategies I and II are ESSs.

Condition 1 of Definition 3.26 tells us that if a^* is an ESS, then the profile $\vec{a}^* = (a^*, a^*)$ is a Nash equilibrium of the fitness game. Indeed, as we have seen in Example 3.9, both (I, I) and (II, II) are Nash equilibria of the Hi-Lo game. However, the concept of ESS does not coincide with the concept of Nash equilibrium. It is not difficult to find symmetric games that have Nash equilibria in pure strategy, but where no strategy is evolutionarily stable. An example is the game of Table 3.1.

Example 3.28. Example 3.2 continued. Consider the fitness game $G = \langle N, X, A, \Phi, \pi \rangle$ corresponding to the game defined in Example 3.2, where the subjective utilities are now interpreted as fitness. We have seen in Example 3.7 that the (fitness) game has two strict Nash equilibria in pure actions: (I, II) and (II, I) . However, when we check if either action is an ESS, we find that:

- $\Phi(\pi(I, I)) = 1 < 5 = \Phi(\pi(II, I))$, hence I is not an ESS;
- $\Phi(\pi(II, II)) = 0 < 2 = \Phi(\pi(I, II))$, hence II is not an ESS.

The inadequacy of the Nash equilibrium concept for an evolutionary analysis emerges even more clearly from the next example.

Example 3.29. Hi-Lo modified. Consider the modified Hi-Lo fitness game depicted in the following table.

	I	II
I	0; 0	0; 0
II	0; 0	2; 2

Action I is *weakly dominated* by action II for both players, because action II yields the same fitness as action I against I , and it yields a strictly higher fitness against II . As in the original Hi-Lo game, however, both (I, I) and (II, II) are Nash equilibria. When we look for ESSs, we find that:

- $\Phi(\pi(II, II)) = 2 > 0 = \Phi(\pi(I, II))$, hence II is an ESS;
- $\Phi(\pi(I, I)) = 0 = \Phi(\pi(II, I))$ and $\Phi(\pi(I, II)) = 0 < 2 = \Phi(\pi(II, II))$, hence I is not an ESS.

Admittedly, we would be surprised to observe a population stably anchored to the Nash equilibrium (I, I) . The rationale behind our surprise is precisely the concept of ESS: as soon as a few II -players arise in the population, the Nash equilibrium (I, I) will be destabilized.

Just as any ESS defines a symmetric Nash equilibrium (by condition 1 of Definition 3.26), in the same way, if (a^*, a^*) is a *strict* Nash equilibrium of the fitness game, then a^* is an ESS (because condition 1 of Definition 3.26 holds strictly by definition of strict Nash equilibrium).

In the absence of ESSs, the static analysis can still appeal to the concept of (*Nash*) *equilibrium state*. Given a fitness game $G = \langle N, X, A, \Phi, \pi \rangle$ with action set $A = \{a^1, \dots, a^m\}$, a *population state* is a probability vector of length m that represents the share of each phenotype in the population. Consequently, we can denote a population state as a mixed action $\alpha = (\alpha(a^1), \dots, \alpha(a^m)) \in \Delta(A)$. At any given state α , each phenotype $a^i \in A$ has an expected fitness defined by

$$\bar{\Phi}(a^i, \alpha) := \sum_{\vec{a}_{-1} \in A^{|N|-1}} \Phi(\pi(a^i, \vec{a}_{-1})) \prod_{j \in N \setminus \{i\}} \alpha(a_j),$$

which, in case of two-player games, reduces to

$$\bar{\Phi}(a^i, \alpha) = \sum_{j=1}^m \Phi(\pi(a^i, a^j)) \alpha(a^j). \quad (3.1)$$

Definition 3.30. Equilibrium state. Given a fitness game G , a population state $\alpha^* \in \Delta(A)$ is a (*Nash*) *equilibrium state* if, for all $a, a' \in A$, it holds that

$$\alpha^*(a) > 0 \Rightarrow \bar{\Phi}(a, \alpha^*) \geq \bar{\Phi}(a', \alpha^*). \quad (3.2)$$

Formula (3.2) expresses the fact that, for a population state to be in equilibrium, all the phenotypes represented in the population must have equal expected fitness. From Definition 3.12 and Lemma 3.15, it is not difficult to see that α^* is an equilibrium state of the fitness game $G = \langle N, X, A, \Phi, \pi \rangle$ if and only if $\bar{\alpha}^*$ is a symmetric Nash equilibrium of the mixed extension $\bar{G} = \langle N, \Delta(X), \Delta(A), \bar{\Phi}, \bar{\pi} \rangle$.

Example 3.31. Example 3.2 continued. Consider the fitness game defined by the game of Table 3.1. As we have seen in Example 3.28, there are no ESSs. However, the polymorphic population state $\alpha^* = (\alpha^*(I) = \frac{1}{3}, \alpha^*(II) = \frac{2}{3})$ is an equilibrium state of the game, since $\bar{\Phi}(I, \alpha^*) = \frac{5}{3} = \bar{\Phi}(II, \alpha^*)$. No other equilibrium states exist for this game. Note that the proportions of phenotypes in the polymorphic equilibrium state correspond to the mixed Nash equilibrium of the game (Example 3.13). In the equilibrium state the mixture of the population keeps all actions equally fit, just as in the mixed Nash equilibrium the mixture of the co-player keeps the player indifferent between her own actions.

Example 3.32. Hi-Lo continued. The Hi-Lo fitness game of Table 3.2 has three equilibrium states: $\alpha^I = (\alpha^I(I) = 1, \alpha^I(II) = 0)$, $\alpha^{II} = (\alpha^{II}(I) = 0, \alpha^{II}(II) = 1)$, and $\alpha^* = (\alpha^*(I) = \frac{2}{3}, \alpha^*(II) = \frac{1}{3})$. The first two correspond to the Nash equilibria in pure actions (Example 3.9), where the monomorphic population is playing an ESS (Example 3.27). The third corresponds to the mixed Nash equilibrium of Example 3.14.

Example 3.32 highlights an important point: whenever an action a^* is an ESS, then the population state α^* such that $\alpha^*(a^*) = 1$ and $\alpha^*(a) = 0$ for all $a \neq a^*$ is an equilibrium state. This follows immediately from Definition 3.26 and Definition 3.30: if $\Phi(\pi(a^*, a^*)) \geq \Phi(\pi(a, a^*))$ for all $a \neq a^*$, then $\bar{\Phi}(a^*, \alpha^*) \geq \bar{\Phi}(a, \alpha^*)$ for all $a \in A$.

The static approach to evolutionary game theory essentially answers two questions: what are the equilibrium states of the population, and what are the monomorphic states which are stable against invasion? There is a third important evolutionary question about the equilibria of a game: what are the *stable* equilibrium states? As we will see shortly, not all the equilibrium states are necessarily stable.

Dynamic analysis Instead of focusing on the static identification of the equilibrium states, the dynamic analysis makes the evolutionary process explicit through the formalization of the dynamics that represent the driving forces within the population, and investigates what equilibrium states are stable with respect to those dynamics. The principal dynamics that have been studied in the literature are two: the *replicator* dynamics (RD) and the *replicator-mutator* dynamics (RMD).

Replicator dynamics The replicator dynamics (see [Taylor and Jonker, 1978], [Hofbauer and Sigmund, 1998]) express the change of proportions of the phenotypes in the population from a generation to the following one when evolution is uniquely determined by (expected) fitness. To present the dynamics, we need to keep track of the changes in the population. The obvious way to proceed is to introduce a time index t for population states $\alpha^0, \alpha^1, \dots, \alpha^t, \alpha^{t+1}$, and so on. The (discrete-time) replicator dynamics for a fitness game G and phenotype $a \in A$ are specified by

$$\alpha^{t+1}(a) = \alpha^t(a) \frac{\bar{\Phi}(a, \alpha^t)}{\bar{\Phi}(\alpha^t)}, \quad (3.3)$$

where $\bar{\Phi}(\alpha^t) := \sum_{a' \in A} \bar{\Phi}(a', \alpha^t) \alpha^t(a')$ is the average fitness in the population at time t . The ratio in equation (3.3) is greater than 1 if the expected fitness of action a is higher than the average fitness of the population. Hence, equation (3.3) formalizes the fact that evolutionary success translates into differential reproduction. If the expected fitness of phenotype a is higher than the average fitness, then the proportion of a -players is expected to raise in the next generation α^{t+1} (vice versa, the proportion of a -players will decrease if a 's expected fitness is lower than the average). Moreover, $\alpha^{t+1}(a) = 0$ for all $a \in A$ such that $\alpha^t(a) = 0$: replicator dynamics only take into account phenotypes that are already present in the population.

Often, in dynamical systems the evolution of the system is described in continuous time. When the population is in state $\alpha \in \Delta(A)$ and time is continuous, the rate of increment (or decrement) of phenotype $a \in A$ is given by:

$$\dot{\alpha}(a) = \alpha(a) (\bar{\Phi}(a, \alpha) - \bar{\Phi}(\alpha)). \quad (3.4)$$

Equation (3.4) states that the proportion of a will be increasing when the fitness of a is higher than the average fitness, so that $\bar{\Phi}(a, \alpha) - \bar{\Phi}(\alpha) > 0$, and decreasing when $\bar{\Phi}(a, \alpha) - \bar{\Phi}(\alpha) < 0$ (as we have seen for the discrete case). As before, we also have that $\dot{\alpha}(a) = 0$ whenever $\alpha(a) = 0$.

Definition 3.33. Rest point. A population state $\alpha^* \in \Delta(A)$ is called a *rest point* of the dynamics if $\dot{\alpha}^*(a) = 0$ for all $a \in A$.

Although discrete-time and continuous-time dynamics do not always coincide (see for example [Weibull, 1995], [Sandholm, 2010], [Cressman and Tao, 2014]),

it is straightforward from equations (3.3) and (3.4) that $\dot{\alpha}^*(a) = 0$ for all $a \in A$ if and only if $\alpha^* = \alpha^{*+1}$. When the population is on a rest point, the proportions of phenotypes are not going to change in the next generations. Furthermore, a population state is in equilibrium when all the phenotypes in the population have the same expected fitness (Definition 3.30). At any equilibrium state α^* it then holds that $\bar{\Phi}(\alpha^*) = \bar{\Phi}(a, \alpha^*)$ for all $a \in A$ with $\alpha^*(a) > 0$. So, if α^* is an equilibrium state it follows from equation (3.4) that $\dot{\alpha}^*(a) = 0$ for all $a \in A$, and from equation (3.3) that $\alpha^* = \alpha^{*+1}$. Any equilibrium state is thus a rest point of the dynamics, but the converse is not true, as shown in the following example.

Example 3.34. Example 3.2 continued. Consider the fitness game defined by Table 3.1, and a population state α^t such that $\alpha^t(I) = \frac{1}{2} = \alpha^t(II)$. Therefore, $\bar{\Phi}(I, \alpha^t) = 1.5$, $\bar{\Phi}(II, \alpha^t) = 2.5$, and $\bar{\Phi}(\alpha^t) = 2$. By computing the discrete-time dynamics, in the next generation we expect $\alpha^{t+1}(I) = \frac{3}{8}$ and $\alpha^{t+1}(II) = \frac{5}{8}$. Looking at the continuous-time RD, we have $\dot{\alpha}^t(I) = -\frac{1}{4}$, and $\dot{\alpha}^t(II) = \frac{1}{4}$. Hence, α^t is not a rest point. We can also check that $\alpha^{t+1}(I) = \frac{1}{2} - \frac{1}{4} \cdot \frac{1}{2} = \frac{3}{8}$, and $\alpha^{t+1}(II) = \frac{1}{2} + \frac{1}{4} \cdot \frac{1}{2} = \frac{5}{8}$, as expected from discrete-time RD. There are three rest points in the game under consideration. Two of them are the monomorphic states. The third is the equilibrium state found in Example 3.31. Notice that the unique equilibrium state of the game is also a rest point, but the two monomorphic rest points are not equilibrium states.

Once some evolutionary dynamics is explicitly defined, a question that naturally arises is: what are the rest points of the dynamic that are evolutionarily stable? In [Maynard Smith, 1982]’s words, “A population is said to be in an evolutionarily stable state if its genetic composition is restored by selection after a disturbance, provided the disturbance is not too large.” More precisely, there are two ways for a rest point to be stable. A weaker sense of stability is captured by the notion of Lyapunov stability.

Definition 3.35. Lyapunov stability. A rest point $\alpha^* \in \Delta(A)$ is said to be *Lyapunov stable* if for any open neighborhood Γ of α^* , there is a neighborhood $\Gamma' \subseteq \Gamma$ of α^* such that

$$\alpha^t \in \Gamma' \Rightarrow \alpha^{t'} \in \Gamma \quad \forall t' > t.$$

Definition 3.35 intuitively expresses that if α^* is Lyapunov stable, then all the population states nearby α^* will stay nearby. A stronger sense in which a state can be evolutionarily stable is given by the following definition.

Definition 3.36. Asymptotic stability. A rest point α^* is called *asymptotically stable* if it is Lyapunov stable and there is an open neighborhood Γ of α^* such that

$$\alpha^t \in \Gamma \Rightarrow \lim_{t \rightarrow \infty} \alpha^t = \alpha^*.$$

A state is asymptotically stable if, besides being Lyapunov stable, there is a neighborhood Γ of α^* where all the points converge to α^* . When this happens, the state α^* is called *attractor*, and the largest neighborhood Γ that has the property of Definition 3.36 is called its *basin of attraction*.

Example 3.37. Example 3.2 continued. We have seen that the fitness game of Table 3.1 has no evolutionarily stable strategy, but it might still have evolutionarily stable states. The candidates are the three rest points that we found in Example 3.34: the polymorphic equilibrium state α^* where $(\alpha^*(I) = \frac{1}{3}, \alpha^*(II) = \frac{2}{3})$, and the two monomorphic states. As for the stability of the polymorphic state α^* , consider the population state α^t with $\alpha^t(I) = \frac{1}{2} = \alpha^t(II)$, as in Example 3.34. We know that in the following generation the population will be in state $\alpha^{t+1}(I) = \frac{3}{8}$ and $\alpha^{t+1}(II) = \frac{5}{8}$. By computing the next time steps we find $\alpha^{t+2} = (\approx 0.3421, \approx 0.6579)$, $\alpha^{t+3} = (\approx 0.3351, \approx 0.6649)$, ..., converging to $\alpha^* = (\frac{1}{3}, \frac{2}{3})$, where the first number of each pair is the proportion of phenotype I and the second number is the proportion of phenotype II . The situation is different when we take into consideration the monomorphic rest points. Let us suppose that the II -monomorphic state is minimally perturbed by a tiny quantity ϵ of I -players that are injected into the population at time t . The expected fitness of phenotype I is then $\Phi(I, \alpha^t) = \epsilon + 2(1 - \epsilon) \approx 2$, whereas the expected fitness of phenotype II is $\Phi(II, \alpha^t) = 5\epsilon + 0(1 - \epsilon) \approx 0$. Consequently, we should expect an increase of I -players in the next generation α^{t+1} . Precisely, I -players will keep increasing until they will occupy $\frac{1}{3}$ of the population. At that moment, their expected fitness will equal that of II -players, and their number will stop growing. By a similar argument, a minimal perturbation of the other monomorphic rest point will result a cascade effect towards the polymorphic equilibrium state α^* . Since any state in the interior of $\Delta(A)$ is attracted to the polymorphic equilibrium state, the basin of attraction of α^* equals the full interior of $\Delta(A)$. The state α^* is thus called a *global* attractor.

It is not difficult to see that being an equilibrium state is a necessary condition for evolutionary stability. If a state α is not in equilibrium, it means that $\alpha(a) > 0$ and $\bar{\Phi}(a, \alpha) < \bar{\Phi}(a', \alpha)$ for some $a, a' \in A$. But then for any open neighborhood Γ of α there is a state $\alpha' \in \Gamma$ with $\alpha'(a') \neq 0$. It follows that it will be possible to find an open neighborhood Γ of α without the property specified in Definition 3.35. The next example shows that being an equilibrium state is not a sufficient condition though: not every equilibrium state is evolutionarily stable.

Example 3.38. Hi-Lo continued. The rest points for the Hi-Lo fitness game are also three, and coincide with the three equilibrium states of Example 3.32: the monomorphic states $\alpha^I = (1, 0)$ and $\alpha^{II} = (0, 1)$, and the polymorphic state $\alpha^* = (\frac{2}{3}, \frac{1}{3})$. The two monomorphic rest points are asymptotically stable, while α^* is not. Obviously, neither α^I nor α^{II} can be a global attractor in this case. It follows that the state space $\Delta(A)$ is split into two basins of attraction. The

frontier of the two basins is exactly the unstable rest point $\alpha^* = (\frac{2}{3}, \frac{1}{3})$: a minimal perturbation towards one of the two phenotypes will start a cascade which will end at the monomorphic state with only that phenotype in the population.

Replicator-mutator dynamics The second dynamic that we want to introduce and make use of in later chapters is an extension of the replicator dynamics, where the proportions of phenotypes in the next generation are determined by the combined effect of fitness-based differential reproduction (replicator dynamics) and possible mutation of some players from one phenotype to another. For this reason, it is called *replicator-mutator* dynamics (see [Hofbauer, 1985], [Nowak, 2006]).

Given a fitness game $G = \langle N, X, A, \Phi, \pi \rangle$ with action set $A = \{a^1, \dots, a^m\}$, the mutator dynamics are expressed by a (right) stochastic matrix

$$M = \begin{pmatrix} M_{11} & \dots & M_{1m} \\ \vdots & \vdots & \vdots \\ M_{m1} & \dots & M_{mm} \end{pmatrix}$$

such that $\sum_{j=1}^m M_{ij} = 1$ for each $i = 1, \dots, m$. The entry M_{ij} fixes the probability of mutation from phenotype a^i to phenotype a^j in the next generation. The discrete-time replicator-mutator dynamics are specified by the joint action of replication and mutation:

$$\alpha^{t+1}(a^i) = \sum_{j=1}^m M_{ji} \cdot \alpha^t(a^j) \frac{\bar{\Phi}(a^j, \alpha^t)}{\bar{\Phi}(\alpha^t)}. \quad (3.5)$$

The formula states that, in each generation t , players play against each other in the population, and each phenotype a^j gets a certain expected fitness that determines its fertility rate. But if some of its offspring are going to mutate into a different phenotype, then the overall share of phenotype a^i in the next generation $t + 1$ is the sum of the expected offspring of each phenotype a^j multiplied by the probability of mutation from a^j into a^i .

In continuous time, RMD are given by:

$$\dot{\alpha}(a^i) = \sum_{j=1}^m \alpha(a^j) \bar{\Phi}(a^j, \alpha) M_{ji} - \alpha(a^i) \bar{\Phi}(\alpha). \quad (3.6)$$

From a biological point of view, it makes sense to postulate that mutation from a phenotype to another happens only very rarely. This amounts to assuming that, for any row i in the stochastic matrix, the probability of mutating $\sum_{j \neq i} M_{ij}$ is substantially lower than the probability of not mutating M_{ii} .

Example 3.39. Example 3.2 continued. Let us consider the fitness game defined by Table 3.1, and suppose the population is at the *II*-monomorphic state

$\alpha^{II} = (0, 1)$. When we study the evolutionary dynamics driven by both replication and mutation, we need to specify the mutation probabilities M_{ij} . Denote by $M_{I,II}$ the probability of mutating from phenotype I to phenotype II , and fix $M_{I,II} = M_{II,I} = \epsilon$, for an arbitrarily small $\epsilon > 0$. Then the population state α^{II} is no longer a rest point of the replicator-mutator dynamics. Indeed, according to RMD, a fraction ϵ of mutants playing action I will arise by mutation in the next generation. This will cause the cascade effect that we have seen in Example 3.37, leading the population to the state $\alpha^* = (\frac{1}{3}, \frac{2}{3})$.

Example 3.40. Hi-Lo continued. Suppose the population is playing the Hi-Lo fitness game, and the population state at time t is $\alpha^t = (\frac{2}{3}, \frac{1}{3})$. As observed in Example 3.38, this is a rest point of the replicator dynamics. If we have $M_{I,II} = M_{II,I} = \epsilon$ again, then it is easy to compute that $\dot{\alpha}^t(I) \neq 0 \neq \dot{\alpha}^t(II)$, so that α^t is not a rest point of RMD in this case. From discrete-time RMD, we get $\alpha^{t+1} = (\frac{2}{3} - \frac{\epsilon}{3}, \frac{1}{3} + \frac{\epsilon}{3})$: we expect an increase of phenotype II and a decrease of phenotype I in the next generation. However, it is still possible for $\alpha^t = (\frac{2}{3}, \frac{1}{3})$ to be a rest point of RMD. If we fix $M_{II,I} = 2M_{I,II}$, for example $M_{I,II} = \epsilon$ and $M_{II,I} = 2\epsilon$, it follows that $\dot{\alpha}^t(I) = \dot{\alpha}^t(II) = 0$. To restore $\alpha^t = (\frac{2}{3}, \frac{1}{3})$ as rest point of the replicator-mutator dynamics in the Hi-Lo game we had to counterbalance the difference in fitness by a proportional difference in mutation rates.

A general result that relates the set of rest points of RMD to the set of rest points of RD is that the limit rest points of RMD constitute a subset of the rest points of RD (see [Samuelson, 1997]), where a *limit* rest point is the limit of the sequence of rest points of RMD as the mutation rates tend to zero.

So far, the evolutionary dynamics have mainly been used to investigate the stability of states, without going much beyond the equilibrium analysis of population games. A second important question that we might ask once the dynamics have been explicitly introduced is about the attainability of different evolutionarily stable states. The attainability of a state is understood as the probability of a random population to evolve to it, and determined by the size of its basin of attraction. In this context, [Huttegger and Zollman, 2010] found examples of games with a unique strict Nash equilibrium and a unique ESS, where nevertheless the population evolved to the ESS less than twenty percent of the times. Yet, ESSs correspond to evolutionarily stable states with respect to most dynamics. This shows that one should be very careful in taking the equilibrium analysis in terms of ESSs as a good approximation for the outcomes of the evolutionary dynamics. An explicit dynamic analysis in terms of basins of attraction and attainability of states is crucial in many cases. We believe that neither a static nor a dynamic analysis is exhaustive by itself, and agree with [Huttegger and Zollman, 2013] on the necessity of a pluralistic approach to the study of evolutionary outcomes. In the following chapters we will therefore make use of both methodologies.

3.3 Population Games With Incomplete Information Under Ambiguity: An Important Example

In this final section we will bring together notions from decision theory and evolutionary game theory by means of an example that is going to introduce us to the evolutionary framework of the following chapters.

Consider a population with two different types of players. Suppose that the members of the population are randomly matched to play the symmetric two-player fitness game $G = \langle \{1, 2\}, X, A, \Phi, \pi \rangle$ depicted in the following table.

	I	II
I	2; 2	1; 0
II	0; 1	5; 5

Here, instead of considering types as different expressed behaviors $a \in A$, we will think of a type as a subjective utility $u : X \rightarrow \mathbb{R}$. This research direction aims at studying the *evolution of preferences* and has been investigated in some recent papers (see, e.g., [Alger and Weibull, 2013], [Robson and Samuelson, 2011], [Dekel et al., 2007]).

Let us denote the set of types in the population by \mathcal{T} , and suppose that there are two different *preference types* $t^1, t^2 \in \mathcal{T}$, with subjective utilities u_{t^1} and u_{t^2} respectively. Furthermore, denote by $\lambda \in [0, 1]$ the proportion of types t^2 , so that the population share of types t^1 will be $1 - \lambda$. The probability of encountering the first type will thus be $P(t^1) = 1 - \lambda$, and analogously $P(t^2) = \lambda$. Given this set-up, when two players are selected from the population and matched to play G , it is natural to model the situation as a Bayesian game with private values $BG = \langle N, S, X, (\Theta_i, T_i, A_i, \tau_i, \vartheta_i, P_i \times P_i, u_i)_{i \in N}, \pi \rangle$ such that:

- $N = \{1, 2\}$
- $T_i = \Theta_i = \mathcal{T} = \{t^1, t^2\}, \forall i \in N$
- $S = \mathcal{T} \times \mathcal{T}$
- $\tau_i : S \rightarrow \mathcal{T}$ is the projection such that $\forall (t^i, t^j) \in \mathcal{T} \times \mathcal{T}, \tau_1(t^i, t^j) = t^i$ and $\tau_2(t^i, t^j) = t^j$
- ϑ_i is the identity function, $\vartheta_i(t^j) = t^j, \forall i, j \in \{1, 2\}$
- X and $A_i (\forall i \in N)$ are as in the fitness game G
- $\pi : \vec{A} \rightarrow X$ is as in the fitness game G , since the game has private values
- $u_i : \Theta_i \times X \rightarrow \mathbb{R}$ is such that $\forall i \in N, \forall t^j \in \mathcal{T}, u_i(t^j, x) = u_{t^j}(x)$




Figure 3.2: The resulting Bayesian population game.

- $P_i \times P_i \in \Delta(\mathcal{T} \times \mathcal{T})$ such that $\forall i \in N, \forall (t^i, t^j) \in \mathcal{T} \times \mathcal{T}, (P_i \times P_i)(t^i, t^j) = P(t^i) \cdot P(t^j)$.

Moreover, suppose that the two preference types in the population have subjective utility u_{t^1} and u_{t^2} such that

1. $u_{t^1} = \Phi$
2. $u_{t^2}(\pi(I, I)) = u_{t^2}(\pi(II, II)) = 0, u_{t^2}(\pi(I, II)) = -4, u_{t^2}(\pi(II, I)) = -2$.

The resulting Bayesian population game is pictured in Figure 3.2. The horizontal dotted line distinguishes worlds where player 1 is of type t^1 (above) from worlds where she is of type t^2 (below), and the vertical dashed line distinguishes worlds where player 2 is of type t^1 (on the left) from those where he is of type t^2 on the right.

However, in realistic situations, it is reasonable to assume that the players might not have access to the precise statistics about the composition of the population, especially when the population is very large and the game has been played only a few times. This amounts to saying that the players cannot pinpoint the actual λ . In such cases, agents may entertain a *non-probabilistic* representation of the uncertainty. This line of thoughts derives from the considerations about the justification of beliefs as a prerequisite for rational choice that we have seen in Chapter 1. Games with incomplete information under ambiguity

(without any evolutionary interpretation though) have recently been studied by [Kajii and Ui, 2005] and [Liu, 2015]. Very much in the spirit of urns à la Ellsberg, it is then natural to think of the random selection of players from the population as an extraction from an urn with unknown composition.

Just as in Ellsberg's examples, whenever the players' uncertainty about the population composition is unmeasurable, a possibility is to represent it as a compact convex set Γ of possible distributions of types in the population. The intuition may be that if the players observed some plays of the game, they could realize that there are some types t^1 , and that there are also some types t^2 , but without being able to narrow down the set of possible proportions of t^2 to more than, for example, $\lambda \in [0.1, 0.9]$. Then, when a player of type t^i is drawn from the population, his or her probability λ of being matched with a type t^2 is within the interval $[0.1, 0.9]$, that specifies a lower probability $\underline{\lambda} = 0.1$ and an upper probability $\bar{\lambda} = 0.9$ of encountering type t^2 .

The unmeasurable uncertainty about the proportions in the population gives rise to a game with incomplete information under ambiguity, that differs from the previous Bayesian games in that the ex ante and interim beliefs of player i are now represented by a *set* of probabilities. Assuming that players have the same uncertainty about the proportions in the population, quantified in terms of a common set Γ of possible distributions of types, it follows that the game has *common* set of priors $\Gamma \otimes \Gamma$, defined as

$$\Gamma \otimes \Gamma := \{P \times P : P \in \Gamma\}$$

where, for all $(t^i, t^j) \in \mathcal{T} \times \mathcal{T}$, $(P \times P)(t^i, t^j) = P(t^i) \cdot P(t^j)$. The last bullet from the previous list, corresponding to the prior beliefs of the players, turns accordingly into:

- $\Gamma \otimes \Gamma \subseteq \Delta(\mathcal{T} \times \mathcal{T})$.

Relying on the convention of denoting simply by $\lambda \in [0, 1]$ the Bernoulli distribution $P_\lambda \in \Gamma$ with parameter λ , we directly express the set Γ as an interval in $[0, 1]$, $\Gamma \subseteq [0, 1]$. Each state $(t^i, t^j) \in \mathcal{T} \times \mathcal{T}$ will then have a lower prior probability $(\underline{P \times P})(t^i, t^j)$ and an upper prior probability $(\overline{P \times P})(t^i, t^j)$. In our example with $\Gamma = [0.1, 0.9]$, we get

$$(\underline{P \times P})(t^1, t^1) = (\underline{P \times P})(t^2, t^2) = 0.01$$

$$(\underline{P \times P})(t^1, t^2) = (\underline{P \times P})(t^2, t^1) = 0.09$$

$$(\overline{P \times P})(t^1, t^1) = (\overline{P \times P})(t^2, t^2) = 0.81$$

$$(\overline{P \times P})(t^1, t^2) = (\overline{P \times P})(t^2, t^1) = 0.25$$

In general, the theoretical literature on decision making under ambiguity has not yet reached a consensus on which updating rule should be used for updating

prior sets of probabilities to posterior sets of probabilities. Note that our formulation implies that players update the prior set of probabilities $\Gamma \otimes \Gamma$ to the posterior set Γ by means of full Bayesian updating (see [Fagin and Halpern, 1990]):

$$\Gamma = \{(P \times P)(\cdot|t^j) \in \Delta(\mathcal{T} \times \mathcal{T}) : P \times P \in \Gamma \otimes \Gamma\}.$$

Games with incomplete information under ambiguity have been introduced and studied in very recent works (see, e.g., [Kajii and Ui, 2005], [Liu, 2015], [Battigalli et al., 2015]). For our goals, it is important to stress that the unmeasurable uncertainty can be resolved to an action choice in different ways.

For brevity, given the Bayesian population game (with common prior $P \times P \in \Gamma \otimes \Gamma$), let us write the ex ante utility for player i and profile $\vec{\sigma}$ as

$$\begin{aligned} & \mathbb{E}_{P \times P}[u_i(\sigma_i, \vec{\sigma}_{-i})] := \\ & \sum_{t_i} (P \times P)(t_i) \sum_{\vec{t}_{-i}} (P \times P)(\vec{t}_{-i}|t_i) u_i(\vartheta_i(t_i), \pi(\vartheta_i(t_i), \vec{\vartheta}_{-i}(\vec{t}_{-i}), \sigma_i(t_i), \vec{\sigma}_{-i}(\vec{t}_{-i}))). \end{aligned}$$

Likewise, we will write the interim utility of player i given profile $\vec{\sigma}$ and type t_i as

$$\begin{aligned} & \mathbb{E}_{P \times P}[u_i(\sigma_i, \vec{\sigma}_{-i})|t_i] := \\ & \sum_{\vec{t}_{-i}} (P \times P)(\vec{t}_{-i}|t_i) u_i(\vartheta_i(t_i), \pi(\vartheta_i(t_i), \vec{\vartheta}_{-i}(\vec{t}_{-i}), \sigma_i(t_i), \vec{\sigma}_{-i}(\vec{t}_{-i}))). \end{aligned}$$

As we have seen in Section 3.1, in Bayesian games a profile of policy functions $\vec{\sigma}^*$ defines a Nash equilibrium in the ex ante strategic form of the game if and only if it defines a Nash equilibrium in the interim strategic form. When uncertainty is unmeasurable and the epistemic state of the players is specified by a set of probability distributions over the states, rather than a single probability, different uncertainty resolution procedures correspond to different equilibrium concepts.

Consider a two-player game with incomplete information and private values like the one above, and a prior set $\Gamma \otimes \Gamma$ of distributions over the states together with a profile $\vec{\sigma} = (\sigma_1, \sigma_2)$. The ex ante “value” of policy σ_i might then be established by looking, for instance, at the minimal expected utility possibly achievable,

$$\min_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \sigma_{3-i})],$$

or, alternatively, by considering the maximal expected utility,

$$\max_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \sigma_{3-i})].$$

The same reasoning holds for the interim utilities. Given a profile $\vec{\sigma}$ and a type t^j , a player may evaluate actions depending on the interim minimal expected utility possibly achievable,

$$\min_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \sigma_{3-i})|t^j]$$

as well as on the interim maximal expected utility possibly achievable,

$$\max_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \sigma_{3-i})|t^j],$$

where, for each $\lambda \in \Gamma$, the expectation is given by:

$$\mathbb{E}_{P_\lambda \times P_\lambda}[u_i(\sigma_i, \sigma_{3-i})|t^j] = (1-\lambda)u_{t^j}(\pi(\sigma_i(t^j), \sigma_{3-i}(t_{3-i}^1))) + \lambda u_{t^j}(\pi(\sigma_i(t^j), \sigma_{3-i}(t_{3-i}^2))).$$

Following the work by [Liu, 2015], we introduce the following equilibrium concepts for the population game under ambiguity.

Definition 3.41. Ex ante Γ -equilibrium. Given an incomplete information game under ambiguity, a profile of policy functions $\vec{\sigma}^*$ is defined to be an *ex ante Γ -equilibrium* if, for all $i \in N$,

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} \min_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \vec{\sigma}_{-i}^*)].$$

The concept of interim Γ -equilibrium is defined analogously.

Definition 3.42. Interim Γ -equilibrium. Given an incomplete information game under ambiguity, a profile of policy functions $\vec{\sigma}^*$ is defined to be an *interim Γ -equilibrium* if, for all $i \in N$ and all $t^j \in \mathcal{T}$,

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} \min_{P \times P \in \Gamma \otimes \Gamma} \mathbb{E}_{P \times P}[u_i(\sigma_i, \vec{\sigma}_{-i}^*)|t^j].$$

Let us now get back to the analysis of the game with incomplete information under ambiguity of our current example. The following table shows the ex ante utilities. Each pair of actions listed in the first column defines a policy $\sigma_1 : T_1 \rightarrow A_1$ for player 1, with the convention that the first element is the action associated with type t^1 and the second element is the action associated with type t^2 (and similarly for player 2). Since the game is symmetric, it suffices to specify the utilities of player 1 only.

	I, I	I, II	II, I	II, II
I, I	$-2\lambda + 2$	$-3\lambda^2 - 3\lambda + 2$	$3\lambda^2 - 4\lambda + 1$	$-5\lambda + 1$
I, II	$-4\lambda + 2$	$3\lambda^2 - 5\lambda + 2$	$-3\lambda^2 + 1$	$-\lambda + 1$
II, I	0	$-9\lambda^2 + 5\lambda$	$\lambda^2 - 10\lambda + 5$	$-9\lambda + 5$
II, II	-2λ	$-3\lambda^2 + 3\lambda$	$3\lambda^2 - 10\lambda + 5$	$-5\lambda + 5$

When $\lambda \in [0.1, 0.9]$, the ex ante minimal expected utilities are those computed in Table 3.3.

On the other hand, the interim analysis of the game is as follows. Given a value of λ , when player 1 observes signal t^1 the interim expected utilities are:

$$\begin{aligned} \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (I, I))|t^1] &= 2 & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (II, I))|t^1] &= 1 + \lambda \\ \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (I, I))|t^1] &= 0 & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (II, I))|t^1] &= 5 - 5\lambda \\ \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (I, II))|t^1] &= 2 - \lambda & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (II, II))|t^1] &= 1 \\ \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (I, II))|t^1] &= 5\lambda & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (II, II))|t^1] &= 5 \end{aligned}$$

	I, I	I, II	II, I	II, II
I, I	0.2; 0.2	-3.13; -1.6	$-\frac{1}{3}(!); 0$	-3.5; -1.8
I, II	-1.6; -3.13	$-\frac{1}{12}(!); -\frac{1}{12}(!)$	-1.43; -2.79	0.1; 0.27
II, I	$0; -\frac{1}{3}(!)$	-2.79; -1.43	-3.19; -3.19	-3.1; -1.57
II, II	-1.8; -3.5	0.27; 0.1	-1.57; -3.1	0.5; 0.5

Table 3.3: Ex ante minimal expected utilities.

When player 1 observes signal t^2 we have instead:

$$\begin{aligned}
\mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (I, I))|t^2] &= 0 & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (II, I))|t^2] &= -4 + 4\lambda \\
\mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (I, I))|t^2] &= -2 & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (II, I))|t^2] &= -2\lambda \\
\mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (I, II))|t^2] &= -4\lambda & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(I, (II, II))|t^2] &= -4 \\
\mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (I, II))|t^2] &= -2 + 2\lambda & \mathbb{E}_{P_\lambda \times P_\lambda}[u_1(II, (II, II))|t^2] &= 0
\end{aligned}$$

Since the game is symmetric the same holds for player 2.

In Table 3.3, we highlighted some expected utilities by means of exclamation marks. The reason of those will remain mysteriously unspecified for the moment, until Section 6.2. For now, the crucial point about this example is that ex ante and interim Γ -equilibria do *not* coincide. Indeed, from the ex ante form in Table 3.3 we can see that the only two ex ante Γ -equilibria are $((I, I), (I, I))$ and $((II, II), (II, II))$. On the contrary, we can check from the interim analysis above that the interim Γ -equilibria are $((I, I), (I, I))$, $((II, II), (II, II))$, and $((I, II), (I, II))$. In general, as shown in [Liu, 2015], and differently than Bayesian games, ex ante and interim Γ -equilibria do no longer coincide in incomplete information games under ambiguity.

Chapter 4

Epistemic Structures for Rationality

Without knowledge action is useless and knowledge without action is futile. (Abu Bakr)

4.1 Introduction

In recent years many game theorists have focused on the epistemic part of playing a game, taking explicitly into account the knowledge and beliefs of the players involved in strategic interactions. The goal of the epistemic approach to game theory is to study what are the players' epistemic conditions that lead to solution concepts. Indeed, there are few games (e.g., the Prisoners' Dilemma) that do not need any strategic reasoning about the others for a rational player to choose an action. In most of the situations instead players have to take into consideration what they think about the others' actions and beliefs in order to choose an action. Furthermore, we need to consider what a player thinks about the others in order to assess her rationality. Strategic thinking comes out when the players reason about the others' actions, knowledge and beliefs, and epistemic game theory makes it explicit.

In the literature there are at least two main formal structures to deal with situations of interactive epistemology: Kripke models, mainly used in logic and computer science [Fagin et al., 1995], and type spaces, more common in economics and game theory ([Harsanyi, 1967], [Harsanyi, 1968a], [Harsanyi, 1968b]). As shown in many papers, both these frameworks can be used to express epistemic conditions for solution concepts. For instance, [Aumann and Brandenburger, 1995] and [Battigalli and Siniscalchi, 2002] state the epistemic conditions for Nash equilibrium and extensive form rationalizability by means of type spaces, whereas in

[Baltag et al., 2009] and [Lorini, 2013] epistemic conditions for backward induction and iterated weak dominance are expressed by Kripke models. The issue we try to address here is to formally study the relationship between these two structures, with a view to a possible broader communication and closer interaction between the two communities, epistemic logic and epistemic game theory.

Quite recently some steps have already been attempted towards the aim of bridging the mentioned frameworks. We are mainly referring to the work in [Heifetz and Mongin, 2001] and in [Zvesper, 2010]. In [Heifetz and Mongin, 2001] the authors are able to identify a logical system which is sound and complete with respect to the class of type spaces based on some modifications of Aumann's system [Aumann, 1995]. Both these logical systems are probabilistic, in the sense that they are expressed in a language with probabilistic operators. Zvesper's work instead starts from a qualitative version of type spaces, that he names *type-space models* and shows that there is an isomorphism with Kripkean state-space models.

In Section 4.2, we start by defining a qualitative multi-agent epistemic language with belief operators. Then, we show firstly how to interpret it on a specific class of qualitative Kripke models, that we call *doxastic game models*, and later we show how to interpret the same language on probabilistic type spaces. Finally, we prove that the two frameworks are semantically equivalent with respect to the language.

In Section 4.3, we extend the language by introducing knowledge operators, in order to express two different epistemic attitudes in our frameworks. We show how to interpret the extended language on type spaces and subsequently we define the corresponding class of *epistemic-doxastic game models*. In the end we prove the semantic equivalence between the two semantics with respect to the extended language. Consequently, *doxastic game models* and *epistemic-doxastic game models* represent the qualitative Kripkean counterpart of type spaces.

Section 4.4 is devoted to the axiomatizations of the logic of belief and of the logic of belief and knowledge interpreted over doxastic game models and epistemic-doxastic game models, respectively. Given the results on the semantic side, these two logics are sound and complete with respect to type spaces too.

We conclude in Section 4.5 by extending the analysis about the equivalence between Kripke-style semantics and type space semantics to a modal language with probabilistic beliefs. The significance of this section consists in showing that the equivalence between the two kinds of semantics is preserved when moving from a qualitative representation of epistemic attitudes to a quantitative one based on probabilities.

4.2 Models With Belief

In this section we will introduce the epistemic structures under consideration: doxastic Kripke models and type spaces. As usual in logic, we start out by introducing a logical language to talk about the epistemic interactive situation of the players in a game, and in the end we interpret it on both doxastic models and type spaces.

4.2.1 Language

Let us be given a static game G , as defined in the previous chapter. We want to endow ourselves with a logical language to talk about the epistemic situation of the players. We define the language $\mathcal{L}_{DGL}(G)$ for the *doxastic game logic* (DGL) as generated by the following grammar:

$$\varphi ::= pl_i(a_i) \mid \neg\varphi \mid \varphi \wedge \varphi \mid B_i\varphi$$

where $a_i \in A_i$ and $i \in N$. The other boolean operators $\vee, \perp, \top, \rightarrow$ and \leftrightarrow are defined in the standard way. The language $\mathcal{L}_{DGL}(G)$ is a doxastic language with a belief operator B_i for each player i . Notice that the language $\mathcal{L}_{DGL}(G)$, and consequently the logic DGL, is parametrized by the game G that we are taking into consideration. It is the game G that gives us the primitives $pl_i(a_i)$ of our language $\mathcal{L}_{DGL}(G)$: for each $a_i \in A_i$ we have one primitive $pl_i(a_i)$, read as “player i is playing her action a_i ”. The doxastic formula $B_i\varphi$ has to be read as “player i believes that φ is true”. Let us define the sets $\Phi_i = \{pl_i(a_i) : a_i \in A_i\}$ and $\Phi = \bigcup_{i \in N} \Phi_i$. Moreover, let us abbreviate $\widehat{B}_i\varphi := \neg B_i\neg\varphi$.

4.2.2 Semantics

4.2.2.1 Doxastic models

In the epistemic logic literature it is standard to express the semantics of doxastic languages by means of structures called *doxastic models*. Doxastic models are a specific type of Kripke models used in modal logic ([Blackburn et al., 2001]).

Definition 4.1 (Doxastic model). A doxastic model is a tuple

$$M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$$

where:

- W is a countable set of possible worlds;
- $v : W \rightarrow 2^\Phi$ is the valuation function for the set of primitives Φ defined in Section 4.2.1;

- $\rightarrow_i \subseteq W \times W$ is the belief relation of player i that satisfies the following conditions:
 - seriality: $\forall w \exists w'$ such that $w \rightarrow_i w'$;
 - transitivity: $\forall u, w, z : u \rightarrow_i w$ and $w \rightarrow_i z$ implies $u \rightarrow_i z$;
 - Euclideaness: $\forall u, w, z : w \rightarrow_i u$ and $w \rightarrow_i z$ implies $u \rightarrow_i z$.

Let us define the belief set of player i at world w as follows: $\rightarrow_i(w) := \{w' : w \rightarrow_i w'\}$.

Before interpreting $\mathcal{L}_{DGL}(G)$ over doxastic models we are going to identify a subclass of doxastic models, that we call *doxastic game models* **DGM** for game G . A similar notion is defined also in [Lorini and Schwarzentruher, 2010].

Definition 4.2 (Doxastic game model). A doxastic game model for the game G is a doxastic model $M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ satisfying the following conditions:

- Adequate valuation condition (AVC):

$$\forall i \in N, \forall w \in W, v_i(w) \text{ is a singleton,}$$

where $v_i(w)$ is the restriction of $v(w)$ to Φ_i , i.e., $v_i(w) = v(w) \cap \Phi_i$;

- Ex interim condition (ExIC):

$$\begin{aligned} \forall i \in N, \forall w, w' \in W, \forall a_i \in A_i, \text{ if } w \rightarrow_i w' \text{ and } pl_i(a_i) \in v(w), \\ \text{then } pl_i(a_i) \in v(w'). \end{aligned}$$

AVC simply says that the valuation function assigns one and only one action to each player at each world, since we do not want to have worlds in which a player can play two different actions at the same time. ExIC means that if a player plays an action, then she believes to play that action. For this reason we call it the *ex interim condition*: it describes a stage in the game where the players have already chosen their own actions, and they might be uncertain only about the others' actions.

Doxastic game models of the form $M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ can be used to provide a semantics for the language $\mathcal{L}_{DGL}(G)$. The following are the truth conditions of formulas in the language $\mathcal{L}_{DGL}(G)$ relative to doxastic game models, where $M, w \models \varphi$ means that formula φ is true at world w in the model M :

- $M, w \models pl_i(a_i)$ iff $pl_i(a_i) \in v(w)$;
- $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$;

- $M, w \models \varphi \wedge \psi$ iff $M, w \models \varphi$ and $M, w \models \psi$;
- $M, w \models B_i \varphi$ iff $\forall w' \in W$, if $w \rightarrow_i w'$ then $M, w' \models \varphi$.

As usual we say that a formula φ is true in a model $M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ if $\forall w \in W, M, w \models \varphi$. Then, a formula φ is valid in **DGM** if φ is true in M for all $M \in \mathbf{DGM}$, and we write $\models_{\mathbf{DGM}} \varphi$. A formula φ is satisfiable in **DGM** if $\neg\varphi$ is not valid in **DGM**.

4.2.2.2 Type spaces

The formal structures mainly used in economics to express epistemic situations in a game are type spaces, introduced in [Harsanyi, 1967], [Harsanyi, 1968a], and [Harsanyi, 1968b]. The classical construction of type spaces is inductive (see [Brandenburger and Dekel, 1993]). Each type t_i of player i is associated with an action $a_i \in A_i = X_i^0$ and represents a hierarchy of beliefs about the other players. The first level of the hierarchy is given by a probability distribution x_i^1 on the actions of the other players $\vec{A}_{-i} = \prod_{j \neq i} A_j = \prod_{j \neq i} X_j^0 = \vec{X}_{-i}^0$, i.e., $x_i^1 \in \Delta(\vec{X}_{-i}^0)$. Call this distribution x_i^1 player i 's 1-order belief, and \vec{X}_{-i}^0 the domain of player i 's 1-order beliefs. Then, anticipating the other players' reasoning, each player i will also form expectations about the other players' beliefs about the other players' action profiles. Incorporating these into the formal account results in a probability distribution $x_i^2 \in \Delta(\vec{X}_{-i}^0 \times \prod_{j \neq i} \Delta(\vec{X}_{-j}^0))$ over the strategy profiles and the first-order beliefs of the other players, where $\vec{X}_{-i}^0 \times \prod_{j \neq i} \Delta(\vec{X}_{-j}^0) = \vec{X}_{-i}^1$ is the domain of player i 's 2-order beliefs. The process goes iteratively: the domain of player i 's $k + 1$ -order beliefs is defined as $\vec{X}_{-i}^k = \vec{X}_{-i}^{k-1} \times \prod_{j \neq i} \Delta(\vec{X}_{-j}^{k-1})$, and a $k + 1$ -order belief of i is a probability distribution $x_i^{k+1} \in \Delta(\vec{X}_{-i}^k)$. Prima facie, if $k > h > 0$ then player i 's k -order belief is more complex than player i 's h -order belief. It was Harsanyi's seminal contribution to see that this process catches itself at infinity, in the sense that player i 's ω -order belief captures player i 's beliefs about the entire hierarchies of beliefs of the other players (see [Brandenburger and Dekel, 1993, Mertens and Zamir, 1985]). That is, a ω -order belief of player i contains information about ω -order beliefs of the other players and not just about lower order beliefs of the other players. This implies that a type of player i is associated with both an action of player i and a probability distribution over the types of the others. Consequently, every type $t_i \in T_i$ of player i can be mapped to an element of A_i and to an element of $\Delta(\vec{T}_{-i})$, with $\vec{T}_{-i} = \prod_{j \neq i} T_j$. Harsanyi's characterization leads to the following simplified definition of type spaces as given by Aumann and Brandenburger [Aumann and Brandenburger, 1995]:

Definition 4.3 (Type space). A type space T for game G is a tuple

$$\langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$$

where:

- T_i is a countable set of types of player i ;
- $\beta_i : T_i \rightarrow \Delta(\vec{T}_{-i})$ is the belief function of player i that associates with each type $t_i \in T_i$ a probability distribution $\mu \in \Delta(\vec{T}_{-i})$ over the types of the others $\vec{T}_{-i} = \prod_{j \neq i} T_j$;
- $\sigma_i : T_i \rightarrow A_i$ is the action function of player i that associates an action $a_i \in A_i$ with each type $t_i \in T_i$.¹

An element $\vec{t} \in \vec{T}$ is called *state*, where as usual $\vec{T} = T_1 \times \dots \times T_n$. Given a state \vec{t} , t_i denotes the element in \vec{t} corresponding to player i .

The truth conditions for formulas of the doxastic language $\mathcal{L}_{DGL}(G)$ relative to a probabilistic type space $T = \langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$ are given by the following clauses. Note that formulas are evaluated at a given state \vec{t} of a type space T :

- $T, \vec{t} \models pl_i(a_i)$ iff $\sigma_i(t_i) = a_i$;
- $T, \vec{t} \models \neg\varphi$ iff $T, \vec{t} \not\models \varphi$;
- $T, \vec{t} \models \varphi \wedge \psi$ iff $T, \vec{t} \models \varphi$ and $T, \vec{t} \models \psi$;
- $T, \vec{t} \models B_i\varphi$ iff $\forall \vec{t}' \in \vec{T}$ if $t'_i = t_i$ and $\beta_i(t_i)(\vec{t}'_{-i}) > 0$ then $T, \vec{t}' \models \varphi$.

The truth condition of the doxastic operator B_i is justified by the way the notion of belief is commonly defined in type spaces. A basic notion in the literature on type spaces is the *event*. An event is a subset $e \subseteq \vec{T}$. An event for player i is a subset $e_{-i} \subseteq \vec{T}_{-i}$. Let $\beta_i(t_i)(e_{-i})$ be the probability that type t_i gives to the event e_{-i} , i.e.,

$$\beta_i(t_i)(e_{-i}) = \sum_{\vec{t}'_{-i} \in e_{-i}} \beta_i(t_i)(\vec{t}'_{-i}).$$

Then, the event that player i believes e_{-i} , denoted by $\mathcal{B}_i(e_{-i})$, is defined in type spaces as follows:

$$\mathcal{B}_i(e_{-i}) = \{\vec{t} \in \vec{T} : \beta_i(t_i)(e_{-i}) = 1\}.$$

Intuitively, $\mathcal{B}_i(e_{-i})$ is the set of states at which player i assigns probability 1 to the event e_{-i} . It is easy to check that $T, \vec{t} \models B_i\varphi$ if and only if $\vec{t} \in \mathcal{B}_i(\{\vec{t}'_{-i} \in \vec{T}_{-i} : T, (t_i, \vec{t}'_{-i}) \models \varphi\})$. In other words, player i believes φ at state \vec{t} if and only if \vec{t}

¹We are considering complete information games: the rules of the game and the subjective utilities are common knowledge among the players. Consequently, unlike [Aumann and Brandenburger, 1995] there is no need to specify a subjective utility function for each type.

belongs to the set of states at which player i assigns probability 1 to set of states at which φ is true.

The notions of validity and satisfiability for formulas of $\mathcal{L}_{DGL}(G)$ relative to the class \mathbf{T} of type spaces are defined in the standard way, as for doxastic game models.

4.2.3 Correspondence

In this section we are going to formally prove what we claimed in the beginning of this chapter, namely that type spaces and doxastic game models are semantically equivalent with respect to the language $\mathcal{L}_{DGL}(G)$. The proof is developed in many steps and it will proceed as follows. Firstly, we present a way to transform type spaces into doxastic models, and in particular it turns out that the resulting doxastic models are doxastic game models. Then, we show that the transformation preserves the truth of all the formulas in $\mathcal{L}_{DGL}(G)$. This gives us the first half of the proof, i.e., the result that if a formula of $\mathcal{L}_{DGL}(G)$ is satisfiable in type spaces, then it is satisfiable in doxastic game models.

The second half of the proof will proceed symmetrically. Firstly we define a way to transform doxastic game models into type spaces, then we prove that the transformation preserves the truth of all the formulas in $\mathcal{L}_{DGL}(G)$. Consequently, we get the second part of the result, i.e., that if a formula of $\mathcal{L}_{DGL}(G)$ is satisfiable in doxastic game models, then it is satisfiable in type spaces.

Starting from a given type space $T = \langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$ we can transform it into a doxastic game model $M^T = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ where:

- $W = \vec{T}$ is the set of worlds (unless differently specified, to ease the notation in what follows we simply write $w \in W$ for the world corresponding to state $\vec{t} \in \vec{T}$, w' for the world corresponding to state \vec{t}' , w'' for the world corresponding to state \vec{t}'' , and so on);
- $v : W \rightarrow 2^\Phi$ is the valuation function, defined such that $pl_i(a_i) \in v(w)$ iff $\sigma_i(t_i) = a_i$;
- \rightarrow_i is the belief relation of player i , defined as follows: $w \rightarrow_i w'$ iff $\beta_i(t_i)(\vec{t}'_{-i}) > 0$ and $t_i = t'_i$.

We have now to prove that any model obtained via this transformation is a doxastic game model. It amounts to showing that any M^T obtained via the transformation satisfies AVC and ExIC.

Proposition 4.4. M^T is a doxastic game model.

Proof. (AVC). Since $\forall i \in N$, $\sigma_i(t_i) = a_i$ for a given $a_i \in A_i$ and $\neg \exists a'_i \in A_i$ s.t. $a'_i \neq a_i$ and $\sigma_i(t_i) = a'_i$, we have that $M^T, w \models pl_i(a_i)$ and $\forall a'_i \in A_i$ s.t. $a'_i \neq a_i$, $M^T, w \not\models pl_i(a'_i)$. Then, for all $w \in W$ $v_i(w)$ is a singleton and AVC is satisfied.

(ExIC). In M^T we have a world $w \in W$ corresponding to each state $\vec{t} \in \vec{T}$. It follows by definition of M^T that if $w' \in \rightarrow_i(w)$ then $t_i = t'_i$. Furthermore, since each type t_i is associated with a unique action $a_i \in A_i$ it follows that if $w \rightarrow_i w'$ and $pl_i(a_i) \in v(w)$ then $pl_i(a_i) \in v(w')$. Hence ExIC is satisfied. \square

One might now wonder why we presented such a transformation among the many possible ones. For instance, a plausible alternative could be to transform the probability distribution into a qualitative ordering \preceq in the following way: $w \preceq_i w'$ iff $\beta_i(t'_i)(\vec{t}'_{-i}) \geq \beta_i(t_i)(\vec{t}_{-i})$ and $t_i = t'_i$. The reason for our choice, as we are going to show in order to conclude the first half of the proof, is that our transformation preserves the truth of all the formulas in our language $\mathcal{L}_{DGL}(G)$. It means that the two structures T and M^T express the same epistemology with respect to the game G taken into account. We can now formally prove this result.

Theorem 4.5. *Let $\vec{t} \in \vec{T}$ be a state in the type space T and $w \in W$ the corresponding world in the doxastic game model M^T built from T . For any φ in $\mathcal{L}_{DGL}(G)$, if $T, \vec{t} \models \varphi$ then $M^T, w \models \varphi$.*

Proof. By induction on the structure of φ . We prove only some cases.

Induction basis: ($\varphi = pl_i(a_i)$). Suppose $T, \vec{t} \models pl_i(a_i)$. Then, by definition of v we have that $M^T, w \models pl_i(a_i)$.

Inductive steps: ($\varphi = B_i\psi$). Suppose $T, \vec{t} \models B_i\psi$. Then, $\forall \vec{t}'$ such that $t'_i = t_i$ and $\beta_i(t_i)(\vec{t}'_{-i}) > 0$, we have $T, \vec{t}' \models \psi$. By inductive hypothesis, $M^T, w' \models \psi$, $\forall w'$ such that $w \rightarrow_i w'$, hence $M^T, w \models B_i\psi$. \square

The second half of the proof consists in the other direction: given an arbitrary doxastic game model M it is always possible to associate with it a corresponding type space T^M . Furthermore, it holds that the doxastic game model M and the associated type space T^M are semantically equivalent with respect to the language $\mathcal{L}_{DGL}(G)$.

Let us be given an arbitrary doxastic game model $M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$. Firstly, we define the types in T^M in the following way: for all $w \in W$ we associate a type t_i of player i such that if $\rightarrow_i(w) = \rightarrow_i(w')$ then $\rightarrow_i(w)$ and $\rightarrow_i(w')$ represent the same type. Formally, $T_i = \{\rightarrow_i(w) : w \in W\}$. Then for any given world w in M we have a state $\vec{t} = (t_1, \dots, t_n)$ in T^M , defined by $\rightarrow_1(w), \dots, \rightarrow_n(w)$. We call \vec{t} the state corresponding to world w . Secondly, we associate with each type t_i an action a_i specified by v : $\sigma_i(t_i) = a_i$ iff $pl_i(a_i) \in v(w)$. By ExIC each type will be associated with a unique action. Finally, we define the probability distribution $\beta_i(t_i)$ over \vec{T}_{-i} by distinguishing two cases: the case in which the support $supp_{t_i}(\beta_i) = \{\vec{t}'_{-i} \in \vec{T}_{-i} : w \rightarrow_i w'\}$ is finite and the case in which $supp_{t_i}(\beta_i)$ is infinite. Let us define first the finite case. For each $i \in N$, for each $t_i \in T_i$ and for each $\vec{t}'_{-i} \in \vec{T}_{-i}$, if $supp_{t_i}$ is finite then:

$$\beta_i(t_i)(\vec{t}'_{-i}) = \begin{cases} \frac{1}{|supp_{t_i}(\beta_i)|} & \text{if } \vec{t}'_{-i} \in supp_{t_i}(\beta_i) \\ 0 & \text{else} \end{cases}$$

Let us now define the infinite case. For each $i \in N$, for each $t_i \in T_i$ and for each $\vec{t}_{-i} \in \vec{T}_{-i}$, if supp_{t_i} is infinite then:

$$\beta_i(t_i)(\vec{t}_{-i}) = \begin{cases} \frac{1}{2^{f(\vec{t}_{-i})}} & \text{if } \vec{t}_{-i} \in \text{supp}_{t_i}(\beta_i) \\ 0 & \text{else} \end{cases}$$

where f is a bijective function with domain $\text{supp}_{t_i}(\beta_i)$ and codomain $\mathbb{N}_+ = \{1, 2, \dots\}$. Since $\text{supp}_{t_i}(\beta_i)$ is countably infinite this function is well-defined and clearly $\sum_{\vec{t}_{-i} \in \text{supp}_{t_i}(\beta_i)} \frac{1}{2^{f(\vec{t}_{-i})}}$ sums up to 1.

To sum up, given an arbitrary doxastic game model $M = \langle W, \rightarrow_1, \dots, \rightarrow_n, v \rangle$, we can associate with it the type space $T^M = \langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$ defined as follows:

- a type t_i of player i for each $\rightarrow_i(w)$: $T_i = \{\rightarrow_i(w) : w \in W\}$;
- an action a_i of t_i specified by v : $\sigma_i(t_i) = a_i$ iff $pl_i(a_i) \in v(w)$;
- a belief function β_i defined as above.

Theorem 4.6. *Let $w \in W$ be a world of M and $\vec{t} \in \vec{T}$ the corresponding state in the type space T^M built from M . For any φ in $\mathcal{L}_{DGL}(G)$, if $M, w \models \varphi$ then $T^M, \vec{t} \models \varphi$.*

Proof. By induction on the structure of φ . We prove only some cases.

Induction basis: ($\varphi = pl_i(a_i)$). Suppose $M, w \models pl_i(a_i)$. Then, by definition of v we have that $T^M, \vec{t} \models pl_i(a_i)$.

Inductive steps: ($\varphi = B_i\psi$). Suppose $M, w \models B_i\psi$. Consequently, $M, w' \models \psi$, $\forall w' \in \rightarrow_i(w)$. By inductive hypothesis, $\forall \vec{t}'$ s.t. $t'_i = t_i$ and $\beta_i(t_i)(\vec{t}'_{-i}) > 0$, $T^M, \vec{t}' \models \psi$, hence $T^M, \vec{t} \models B_i\psi$. \square

The following corollary finally states that the class of type spaces \mathbf{T} and the class of doxastic game models \mathbf{DGM} provide equivalent semantics with respect to the language $\mathcal{L}_{DGL}(G)$.

Corollary 4.7. *A formula φ of $\mathcal{L}_{DGL}(G)$ is satisfiable in \mathbf{T} iff it is satisfiable in \mathbf{DGM} .²*

²In order to avoid further complications, here we limit ourselves to the countable case: countable type sets and countable possible worlds set. However, our results generalize to the uncountable case by appropriately endowing type sets T_i and \vec{T}_{-i} with σ -algebras, and by ensuring that every $\beta_i(t_i)$ is a measurable function, as usual in type spaces literature (see [Aumann and Brandenburger, 1995] Section 6).

4.3 Models With Belief and Knowledge

In this section we want to extend our language to deal with more than one epistemic attitude. A straightforward extension is to introduce another operator to represent “knowledge” or “absolute certainty with no possibility at all for error”. We want to show if and how we can interpret it on our structures, proving the same equivalence result between type spaces and Kripke models with respect to the enriched language.

4.3.1 Language

Given a static game G , we define an extension of our logic DGL, and we call it *epistemic-doxastic game logic* EDGL. The language $\mathcal{L}_{EDGL}(G)$ of EDGL is defined by the following grammar:

$$\varphi ::= pl_i(a_i) \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid B_i\varphi \mid \Box\varphi$$

where $a_i \in A_i$ and $i \in N$.

The language $\mathcal{L}_{EDGL}(G)$ is the doxastic game language $\mathcal{L}_{DGL}(G)$ of DGL extended by the knowledge operator K_i , and the universal operator \Box , which turns out to be useful for obtaining the axiomatization, as we will see in Section 4.4. $K_i\varphi$ is read as “player i knows that φ is true”, while $\Box\varphi$ is read as “ φ is universally true”. As before, let us abbreviate the dual of K_i as $\widehat{K}_i\varphi := \neg K_i\neg\varphi$ and the dual of \Box as $\Diamond\varphi := \neg\Box\neg\varphi$.

4.3.2 Semantics

4.3.2.1 Type spaces

In a similar way to what we did above we want to interpret our language $\mathcal{L}_{EDGL}(G)$ over a type space T as semantics. In order to do that, we simply need to add to the previous list of clauses of Section 4.2.2.2 the following clauses for K_i -formulas and \Box -formulas:

- $T, \vec{t} \models K_i\varphi$ iff $\forall \vec{t}' \in \vec{T}$ if $t'_i = t_i$ then $T, \vec{t}' \models \varphi$;
- $T, \vec{t} \models \Box\varphi$ iff $\forall \vec{t}' \in \vec{T}$, $T, \vec{t}' \models \varphi$.

Looking at these clauses, we can observe that the K_i -operator ranges over all the states with the same type for player i , whereas the \Box -operator ranges over all the states in \vec{T} . We will spend more words in Section 4.3.5 on the interpretation. In particular, although the clause for K_i -formulas seems the obvious one for type spaces, we will see that it is far from being uncontroversial and it will have important consequences on Kripke models.

4.3.2.2 Epistemic-doxastic models

In the literature about Kripke models, the semantics usually associated with epistemic-doxastic languages containing both a B_i -operator and a K_i -operator are called *epistemic-doxastic models*, or just *epistemic models* for brevity (see [Kraus and Lehmann, 1988]). Epistemic models are nothing but multi-relational Kripke models commonly used in modal logic ([Blackburn et al., 2001]).

Definition 4.8 (Epistemic-doxastic model). An epistemic-doxastic model is a tuple $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ where:

- W is a countable set of possible worlds;
- $v : W \rightarrow 2^\Phi$ is the valuation function for the set Φ defined in Section 4.2.1;
- \sim_i is the epistemic accessibility relation of player i , that is an equivalence relation over W ;
- $\rightarrow_i \subseteq W \times W$ is the belief relation of player i which satisfies the following conditions:
 - seriality: $\forall w \exists w'$ s.t. $w \rightarrow_i w'$;
 - $\rightarrow_i \subseteq \sim_i$;
 - $\forall w, w' \in W$, if $w \sim_i w'$ then $\rightarrow_i(w) \subseteq \rightarrow_i(w')$.

Moreover, let us write $\sim_i(w)$ for the partition cell (also called *information set*) of player i containing world w : $\sim_i(w) = \{w' \in W : w \sim_i w'\}$.

In order to interpret $\mathcal{L}_{EDGL}(G)$ over epistemic models and to state a result of semantic equivalence with respect to type spaces, we are going to identify a subclass of epistemic models, that we call *epistemic-doxastic game models*, or just *epistemic game models* **EGM**.

Definition 4.9 (Epistemic game model). Epistemic game models are epistemic-doxastic models satisfying the conditions AVC and ExIC given in Definition 4.2, and the following condition:

- Epistemic independence condition (EIC):

$$\sim_1(w_1) \cap \dots \cap \sim_n(w_n) \neq \emptyset \text{ for every } (w_1, \dots, w_n) \in W^n.$$

Roughly speaking, EIC says that each player has no reason to rule out any possible information set of the others specified in the model: if an information set of i is present in the model, then the other players should not consider it impossible at any world.

Epistemic game models can be used to represent a semantics for the language $\mathcal{L}_{EDGL}(G)$. The semantic clauses are the same as for doxastic game models, plus the following clauses for K_i -formulas and \Box -formulas:

- $M, w \models K_i \varphi$ iff $\forall w' \in W$ if $w \sim_i w'$ then $M, w' \models \varphi$;
- $M, w \models \Box \varphi$ iff $\forall w' \in W, M, w' \models \varphi$.

As for **DGM**, we say that a formula φ is true in a model M if $\forall w \in W, M, w \models \varphi$. Then, a formula φ is valid in **EGM** (and we write $\models_{EGM} \varphi$) if φ is true in M for all $M \in \mathbf{EGM}$, and a formula φ is satisfiable in **EGM** if $\neg\varphi$ is not valid in **EGM**.

We want to point out also that the condition ExIC in epistemic game models implies that each player knows her own action. This makes sense, since we are describing an ex interim stage of the game, where each player is already certain about her own choice. This is expressed by the following validity:

Lemma 4.10. $\forall i \in N, \forall a_i \in A_i, \models_{EGM} pl_i(a_i) \rightarrow K_i pl_i(a_i)$.

Proof. By contradiction, suppose that $M, w \models pl_i(a_i)$ and that $M, w \not\models K_i pl_i(a_i)$. Consequently, $\exists u \in \sim_i(w)$ s.t. $M, u \not\models pl_i(a_i)$. Hence, by AVC $\exists a'_i \in A_i$ s.t. $a'_i \neq a_i$ and $M, u \models pl_i(a'_i)$. By definition of \rightarrow_i we have that if $w \sim_i u$ then $\rightarrow_i(w) = \rightarrow_i(u)$. It follows that $\forall w' \in W$, if $w' \in \rightarrow_i(w)$ then $w' \in \rightarrow_i(u)$ too. Then, by ExIC $M, w' \models pl_i(a_i)$ and $M, w' \models pl_i(a'_i)$. Contradiction with AVC. \square

4.3.3 Correspondence

In this section we prove the semantic equivalence between type spaces and epistemic game models with respect to the language $\mathcal{L}_{EDGL}(G)$. The proof will proceed in the same way as before, i.e., it will be divided into two parts and we will make use of a transformation of one structure into the other in order to show the equivalence.

Given an arbitrary type space T , the corresponding epistemic-doxastic model M^T is defined as the tuple $M^T = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ where:

- $W = \vec{T}$ is the set of worlds;
- $v : W \rightarrow 2^\Phi$ is the valuation function such that $pl_i(a_i) \in v(w)$ iff $\sigma_i(t_i) = a_i$;
- \sim_i is the accessibility relation of player i , given by: $\forall w, w' \in W, w \sim_i w'$ iff $t_i = t'_i$. Then \sim_i determines a partition over W ;
- \rightarrow_i is the belief relation of player i , defined as follows: $w \rightarrow_i w'$ iff $\beta_i(t_i)(\vec{t}'_{-i}) > 0$ and $t_i = t'_i$.

We have now to show that the epistemic-doxastic model we obtain via the transformation is an epistemic game model, namely we have to show that it satisfies AVC, ExIC and EIC.

Proposition 4.11. M^T is an epistemic game model.

Proof. The proof for AVC and ExIC is similar to Proposition 1. Here we only prove that EIC holds too.

(EIC). Since $\vec{T} = T_1 \times \dots \times T_n$, each state $\vec{t} \in \vec{T}$ has the form (t_1, \dots, t_n) , and $\forall t_i \forall \vec{t}_{-i} \exists \vec{t}$ s.t. $\vec{t} = (t_i, \vec{t}_{-i})$. Denoting by $\sim_i(t_i)$ the partition cell corresponding to type t_i , it follows that $\forall (t_1, \dots, t_n) \in \vec{T} \exists w \in W$ s.t. $\sim_1(t_1) \cap \dots \cap \sim_n(t_n) = w$. Consequently, $\sim_1(w_1) \cap \dots \cap \sim_n(w_n) \neq \emptyset$ for every $(w_1, \dots, w_n) \in W^n$ and EIC is satisfied. \square

To conclude the first half of the proof we show that this transformation into epistemic game models preserves the truth of all the formulas of $\mathcal{L}_{EDGL}(G)$. The proof is similar to the previous one, so in this section we can just focus on the part for K_i -formulas and \Box -formulas.

Theorem 4.12. *Let $\vec{t} \in \vec{T}$ be a state in T and let $w \in W$ be the corresponding world in M^T . For any φ in $\mathcal{L}_{EDGL}(G)$, if $T, \vec{t} \models \varphi$ then $M^T, w \models \varphi$.*

Proof. By induction on the structure of φ .

($\varphi = K_i\psi$). Suppose $T, \vec{t} \models K_i\psi$. Then, $\forall \vec{t}'$ s.t. $t'_i = t_i$ $T, \vec{t}' \models \psi$. By inductive hypothesis, $M^T, w' \models \psi$, $\forall w' \in \sim_i(w)$, hence $M^T, w \models K_i\psi$.

($\varphi = \Box\psi$). Suppose $T, \vec{t} \models \Box\psi$. Then, $\forall \vec{t}' \in \vec{T}$, $T, \vec{t}' \models \psi$. By inductive hypothesis, $M^T, w' \models \psi$, $\forall w' \in W$, hence $M^T, w \models \Box\psi$. \square

In the second part of the proof we are going to show the other direction: given an arbitrary epistemic game model M it is always possible to associate with it a corresponding type space T^M . Moreover, it holds again that the epistemic game model M and the associated type space T^M are semantically equivalent with respect to the language $\mathcal{L}_{EDGL}(G)$.

Let us be given an arbitrary epistemic game model M . Firstly, we define a type $t_i \in T_i$ for each i 's partition cell $\sim_i(w)$ in M : $T_i = \{\sim_i(w) : w \in W\}$. Notice that when we have \sim_i relations in the model this is equivalent to the definition of types given for Theorem 4.6. Then by EIC, given n arbitrary partition cells, one for each player, the intersection will always be non-empty. The worlds w in the intersection are associated with the state $\vec{t} = (t_1, \dots, t_n)$, where types t_1, \dots, t_n are determined by the partition cells $\sim_i(w)$ for all i : let us call \vec{t} the state corresponding to those worlds w in the intersection. Moreover, EIC guarantees that in T^M the states \vec{T} correspond to the Cartesian product of type sets T_i , i.e. $\vec{T} = T_1 \times \dots \times T_n$. Secondly, we associate with each type t_i an action a_i specified by v : $\sigma_i(t_i) = a_i$ iff $pl_i(a_i) \in v(w)$. By ExIC each type will be associated with a unique action. Finally, for each type t_i we define the probability distribution $\beta_i(t_i)$ over \vec{T}_{-i} in the same way as we did in the proof of Theorem 4.6 by distinguishing two cases, the case in which the support $supp_{t_i}(\beta_i) = \{\vec{t}_{-i} \in \vec{T}_{-i} : w \rightarrow_i w'\}$ is finite and the case in which it is infinite, with w being the world in W associated with the type t_i .

To sum up, taken an arbitrary $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$, we can associate with it the corresponding type space $T^M = \langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$ defined as follows:

- a type t_i of player i for each i 's partition cell $\sim_i(w)$ in M : $T_i = \{\sim_i(w) : w \in W\}$;
- an action a_i of t_i specified by v : $\sigma_i(t_i) = a_i$ iff $pl_i(a_i) \in v(w)$;
- a belief function β_i defined differently for the finite case and for the infinite case, as before.

Theorem 4.13. *Let $w \in W$ be a world of M and $\vec{t} \in \vec{T}$ the corresponding state in T^M . For any φ in $\mathcal{L}_{EDGL}(G)$, if $M, w \models \varphi$ then $T^M, \vec{t} \models \varphi$.*

Proof. By induction on the structure of φ .

($\varphi = K_i\psi$). Suppose $M, w \models K_i\psi$. Then, by definition $M, w' \models \psi$, $\forall w' \in \sim_i(w)$. By inductive hypothesis, $\forall \vec{t}'$ s.t. $t'_i = t_i$, $T^M, \vec{t}' \models \psi$, hence $T, \vec{t} \models K_i\psi$.

($\varphi = \Box\psi$). Suppose $M, w \models \Box\psi$. Then, by definition $M, w' \models \psi$, $\forall w' \in W$. By inductive hypothesis, $\forall \vec{t}' \in \vec{T}$, $T^M, \vec{t}' \models \psi$, hence $T^M, \vec{t} \models \Box\psi$. \square

The following corollary then states that the class of type spaces \mathbf{T} and the class of epistemic game models \mathbf{EGM} provide equivalent semantics with respect to the language $\mathcal{L}_{EDGL}(G)$.

Corollary 4.14. *A formula φ of $\mathcal{L}_{EDGL}(G)$ is satisfiable in \mathbf{T} iff it is satisfiable in \mathbf{EGM} .*

4.3.4 Examples

We consider a two player game G , where Ann has action set $\{U, D\}$ and Bob has action set $\{L, R\}$. Then, an example of type space for G is pictured in Figure 4.1.

Figures 4.2 and 4.3 depict the transformation into epistemic game models as defined above. Arrows represent the belief relation: an arrow going from w to w' means $w \rightarrow_i w'$. Squared boxes that partition the set of all possible worlds represent the accessibility relation and each box corresponds to a type.

Notice that if we drop the ex interim condition we can represent a situation in which players have not decided yet their actions and they do not know what their own action will be (Figure 4.4).

		Bob			
		$1L$	$2L$	$1R$	$2R$
Ann	$1U$	$\frac{1}{2}, \frac{1}{3}$	$\frac{1}{2}, 1$	$0, \frac{1}{3}$	$0, \frac{1}{4}$
	$2U$	$\frac{1}{3}, 0$	$\frac{2}{3}, 0$	$0, 0$	$0, 0$
	$1D$	$\frac{1}{2}, 0$	$\frac{1}{2}, 0$	$0, 0$	$0, \frac{3}{4}$
	$2D$	$\frac{1}{2}, \frac{2}{3}$	$0, 0$	$\frac{1}{2}, \frac{2}{3}$	$0, 0$

Figure 4.1: A possible type space




Figure 4.2: Transformation: Ann

4.3.5 Discussion

After having presented the equivalence results between the Kripkean semantics and type space semantics for the qualitative epistemic languages of belief and knowledge, we want to talk about two conceptual issues that are relevant here. First of all, we spend some more words on the concept of knowledge introduced in Section 4.3. Secondly, we briefly discuss the distinction between probabilistic (quantitative) type spaces and qualitative type spaces.




Figure 4.3: Transformation: Bob

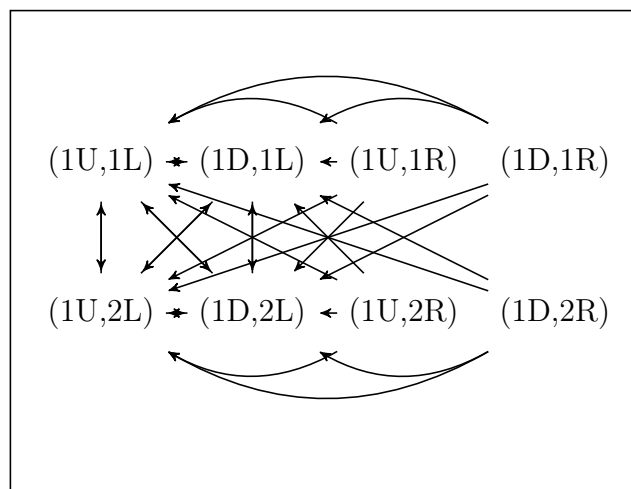


Figure 4.4: Transformation without ExIC: Ann type 1

4.3.5.1 On the knowledge operator

One might wonder why we introduced the operator K_i . In fact, only one basic epistemic operator is normally introduced and used in type spaces: the probability 1-operator, that corresponds to the operator B_i . In economic literature this probability 1-operator is sometimes called knowledge and sometimes belief (see [Aumann and Brandenburger, 1995] and [Brandenburger, 2008]). Indeed,

Aumann and Brandenburger [Aumann and Brandenburger, 1995] write:

In this paper, “know” means “ascribe probability 1 to”. This is sometimes called “believe”, while “know” is reserved for absolute certainty with no possibility at all for error.

The reason why we introduced two different epistemic operators is related to that observation. As mentioned in the previous quote, in (modal) epistemic literature there are two different basic notions, knowledge and belief, that standardly correspond to a S5-operator and a KD45-operator. Then, it makes sense to try to express these two different epistemic attitudes in both the structures we are dealing with, and to compare their interpretations in parallel.

The most obvious thing to do if we want to represent in type spaces both mentioned epistemic attitudes, absolute certainty with no possibility for error and probability 1 (or, knowledge and belief), is to interpret the knowledge operator as we showed in Section 4.3.2.1. Far from being uncontroversial, this definition of the knowledge operator on type spaces gives rise to interesting properties that make the K_i -operator different from the classical S5-knowledge operator used in logic, computer science and distributed artificial intelligence [Fagin et al., 1995]. As we will see, the main difference from the standard knowledge operator is made by the condition EIC. The following propositions show two properties satisfied by the K_i -operator but not by the classical S5-knowledge operator.

Proposition 4.15. *For all $i, j \in N$ such that $i \neq j$, we have $\models_{EGM} K_i K_j \varphi \leftrightarrow \Box \varphi$.*

Proof. (right-to-left). Let $i \neq j$. Suppose that $M, w \models K_i K_j \varphi$. Then, $\forall w' \in \sim_i(w)$, $M, w' \models K_j \varphi$. However, every world $u \in W$ belongs to a partition cell of player j , i.e., $\sim_j(u)$, and, by EIC, for every partition cell $\sim_j(u)$ of player j we have that $\sim_j(u) \cap \sim_i(w) \neq \emptyset$. Then, in every partition cell $\sim_j(u)$ of player j there is a world w' s.t. $M, w' \models K_j \varphi$. It means that $\forall u \in W$, $M, u \models \varphi$. Hence, $M, w \models \Box \varphi$.

(left-to-right). Trivial. □

Proposition 4.16. *For all $i \in N$, we have $\models_{EGM} \Diamond pl_{-i}(\vec{a}_{-i}) \leftrightarrow \widehat{K}_i pl_{-i}(\vec{a}_{-i})$.*

Proof. (left-to-right). Suppose that $M, w \models \Diamond pl_{-i}(\vec{a}_{-i})$. Then, $\exists w' \in W$ s.t. $M, w' \models pl_{-i}(\vec{a}_{-i})$, with $pl_{-i}(\vec{a}_{-i}) := \prod_{j \neq i} pl_j(a_j)$. For each $j \in N$ s.t. $j \neq i$ and for each $pl_j(a_j) \in pl_{-i}(\vec{a}_{-i})$, by ExIC there must be a partition cell $\sim_j(w')$ where $\forall u \in \sim_j(w')$, $M, u \models pl_j(a_j)$. By EIC, we have that: $\forall u' \in W$, $\sim_i(u') \cap \bigcap_{j \neq i} \sim_j(w') \neq \emptyset$. Hence, $M, w \models \widehat{K}_i pl_{-i}(\vec{a}_{-i})$.

(right-to-left). Trivial. □

Proposition 4.16 says that if a particular action profile $pl_{-i}(\vec{a}_{-i})$ is possible, i.e., if $pl_{-i}(\vec{a}_{-i})$ holds at some world in the model, then player i knows for sure that it is possible. Proposition 4.15 on the other hand is a more general and

stronger property of the model: it states that if i and j are different players, then player i knows that player j knows that φ if and only if φ is necessary.

Let us spend some words on the meaning of Proposition 4.15. As it is clear from the definition of type space given in Section 4.2.2.2, a type for player i provides information about the ‘psychological situation’ of player i , namely: (i) her actual choice, and (ii) her subjective probability distribution over possible states. As emphasized above, the operator K_i should be interpreted as an operator of absolute unrevisable certainty, in the sense that $K_i\varphi$ is true if and only if φ is true in all states that player i envisages (or imagines) as possible. The truth condition of this operator in the type space semantics given in Section 4.3.2.1 presupposes that for every possible psychological situation of the other players, player i envisages a state in which this psychological situation occurs. In other words, player i only excludes from her information set those states in which her psychological situation is different from her actual psychological situation. Under this assumption, Proposition 4.15 makes perfect sense. Indeed, let i and j be different players. The previous assumption implies that, for every possible state $\vec{t}'' \in \vec{T}$ that is not in the information set of player i at the actual state \vec{t} , there exists a state $\vec{t}' \in \vec{T}$ in player i ’s information set at the actual state \vec{t} such that \vec{t}'' is included in player j ’s information set at state \vec{t}' . Therefore, clearly, if $K_iK_j\varphi$ is true at state \vec{t} then $\Box\varphi$ is true at \vec{t} too. The other direction of the equivalence (i.e., $\Box\varphi$ implies $K_iK_j\varphi$) holds for obvious reasons.

4.3.5.2 Qualitative vs. quantitative type spaces

In Section 4.2.2.2 we have introduced probabilistic type spaces as defined by Aumann and Brandenburger [Aumann and Brandenburger, 1995] and justified this definition on the basis of Harsanyi’s characterization. It is worth noting that Harsanyi’s characterization only holds because of the properties of probabilities. In particular, in Harsanyi’s type spaces probabilities are σ -additive, and therefore continuous on increasing and decreasing sequences of events. As it has been shown by [Fagin et al., 1999] (see also [Fagin, 1994, Fagin et al., 1991, Heifetz and Samet, 1998]), there is an analogous inductive construction of *qualitative* type spaces that does not satisfy the property that all information about other players’ beliefs is captured at level ω , as the construction might need to carry out transfinitely long. Type spaces studied by Fagin et al. are qualitative: given player i ’s basic domain of uncertainty $W_i^0 = \vec{A}_{-i}$, f_0 denotes a member of W_i^0 . Each assignment $\langle f_0 \rangle$ represents a “possible 1-world”, and the domain of i ’s 1-order beliefs is the set of all possible 1-worlds W_i^1 . Then, i ’s 1-order belief is defined as a set $f_1(i) \subseteq W_i^1$. Inductively, player i ’s k -level belief is defined as a set $f_k(i) \subseteq W_i^k$ of possible k -worlds, i.e., k -tuples of functions $\langle f_0, \dots, f_{k-1} \rangle$. Equivalently, a 1-order belief of i can be expressed as a function $f_1(i) : W_i^1 \rightarrow \{0, 1\}$, and i ’s k -order belief as a function $f_k(i) : W_i^1 \times \dots \times W_i^k \rightarrow \{0, 1\}$. Specifically, it is shown by Fagin et al. that for qualitative type spaces it is not necessarily the

case that a ω -order belief $f_\omega(i)$ of i contains information about ω -order beliefs of the other players. In this sense, there is no simplified definition of qualitative type spaces which is analogous to Definition 4.3 in Section 4.2.2.2 for quantitative type spaces and which is justified in the light of an argument à la Harsanyi. However, since we do not use such an inductive construction of qualitative hierarchies of beliefs and we deal with finitary modal logics, the results by Fagin et al. are not problematic from our standpoint.

The reason why we have introduced probabilistic (quantitative) type spaces instead of qualitative type spaces à la Fagin et al. is that in Section 4.5 we will move from a qualitative representation of epistemic attitudes to a quantitative representation which requires a probabilistic interpretation in terms of probabilistic type spaces. Thus, we preferred to state semantic equivalence with respect to a unique type space representation that applies both to qualitative and quantitative languages. However, it is worth noting that, as for the interpretation of the qualitative epistemic language with belief operators given in Section 4.2 and of its extension by knowledge operators given in Section 4.3, we could have completely omitted the probabilistic aspect of type spaces as defined in Definition 4.3 and given an analogous qualitative definition in which functions β_i are replaced by functions $\beta'_i : T_i \rightarrow 2^{\vec{T}^{-i}}$, where $\vec{t}_{-i} \in \beta'_i(t_i)$ means that player i 's type t_i considers type \vec{t}_{-i} possible and, viceversa, $\vec{t}_{-i} \notin \beta'_i(t_i)$ means that player i 's type t_i considers type \vec{t}_{-i} impossible. The two semantics, the one with functions β_i of Definition 4.3 and the one in which functions β_i are replaced by functions β'_i , have clearly the same sets of validities for both qualitative epistemic languages introduced here. We conjecture that the sets of validities for the two qualitative epistemic languages do not change if we adopt the qualitative type space semantics à la Fagin et al. in which type spaces are defined in an inductive way. We postpone the proof of this conjecture to future work.

4.4 Axiomatization

In this section we provide sound and complete axiomatizations for the logics EDGL and DGL relative to the class of epistemic game models and doxastic game models, respectively.

Given the equivalences between epistemic game models and type spaces with respect to EDGL (Corollary 4.29) and between doxastic game models and type spaces with respect to DGL (Corollary 5.5), these axiomatizations will also turn out to be sound and complete relative to type spaces.

Theorem 4.17. *The set of validities of the logic EDGL relative to the class of epistemic game models (**EGM**) is completely axiomatized by the principles given in Figure 4.5.*

Proof. Proving that the axioms given in Figure 4.5 are sound with respect to the

• **Axioms for EDGL:**

- (1) All tautologies of classical propositional logic
- (2) Axioms K, T, 4 and B for the universal modality \Box
 - (a) $\Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$
 - (b) $\Box\varphi \rightarrow \varphi$
 - (c) $\Box\varphi \rightarrow \Box\Box\varphi$
 - (d) $\varphi \rightarrow \Box\Diamond\varphi$
- (3) Axioms K, T, 4 and B for the knowledge modality K_i
 - (a) $K_i(\varphi \rightarrow \psi) \rightarrow (K_i\varphi \rightarrow K_i\psi)$
 - (b) $K_i\varphi \rightarrow \varphi$
 - (c) $K_i\varphi \rightarrow K_iK_i\varphi$
 - (d) $\varphi \rightarrow K_i\widehat{K}_i\varphi$
- (4) Axioms K and D for the belief modality B_i
 - (a) $B_i(\varphi \rightarrow \psi) \rightarrow (B_i\varphi \rightarrow B_i\psi)$
 - (b) $\neg(B_i\varphi \wedge B_i\neg\varphi)$
- (5) Interaction axioms between universal modality, knowledge modality and belief modality
 - (a) $\Box\varphi \rightarrow K_i\varphi$
 - (b) $K_i\varphi \rightarrow B_i\varphi$
 - (c) $B_i\varphi \rightarrow K_iB_i\varphi$
 - (d) $(\Diamond K_1\varphi_1 \wedge \dots \wedge \Diamond K_n\varphi_n) \rightarrow \Diamond(K_1\varphi_1 \wedge \dots \wedge K_n\varphi_n)$
- (6) Axioms for the atomic formulas $pl_i(a_i)$
 - (a) $\bigvee_{a_i \in A_i} pl_i(a_i)$
 - (b) $pl_i(a_i) \rightarrow \neg pl_i(a'_i)$ if $a_i \neq a'_i$
 - (c) $pl_i(a_i) \rightarrow K_i pl_i(a_i)$

• **Rules of inference for EDGL:**

- (7) From φ and $\varphi \rightarrow \psi$ infer ψ
- (8) From φ infer $\Box\varphi$

Figure 4.5: Axiomatization of EDGL

class **EGM** and that the inference rules preserve validity is just a routine task and we do not give it here.

As to completeness, let us define the class of *weak* epistemic game models

(**WEGM**) as the class of epistemic models that satisfy the epistemic independence condition (EIC) but do not necessarily satisfy the adequate valuation condition (AVC) and the ex interim condition (ExIC). In other words, epistemic game models are a subclass of *weak* epistemic game models that satisfy both the adequate valuation condition (AVC) and the ex interim condition (ExIC).

We write $\models_{WEGM} \varphi$ to mean that the EDGL-formula φ is *valid* relative to the class **WEGM**.

Moreover, for any finite set Δ of EDGL-formulas, we write $\Delta \models_{WEGM} \varphi$ to mean that φ is a *logical consequence* of the set of formulas Δ relative to the class **WEGM**. That is, $\Delta \models_{WEGM} \varphi$ iff, for every *weak* epistemic game model $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$, if $M, w \models \bigwedge_{\psi \in \Delta} \psi$ for all $w \in W$, then $M, w \models \varphi$ for all $w \in W$.

The following Proposition 4.18 highlights that the validity problem relative to the class **EGM** is reducible to the logical consequence problem relative to the class **WEGM**.

Proposition 4.18. *Let*

$$\begin{aligned} \Delta_0 = & \{ \bigvee_{a_i \in A_i} pl_i(a_i) : i \in N \} \cup \\ & \{ pl_i(a_i) \rightarrow \neg pl_i(a'_i) : i \in N \text{ and } a_i, a'_i \in A_i \text{ with } a_i \neq a'_i \} \cup \\ & \{ pl_i(a_i) \rightarrow B_i pl_i(a_i) : i \in N \text{ and } a_i \in A_i \} \end{aligned}$$

Then, for every EDGL-formula φ , $\models_{EGM} \varphi$ iff $\Delta_0 \models_{WEGM} \varphi$.

Proof. We just need to observe that the (global) axioms in Δ_0 force a *weak* epistemic game model to satisfy the ex interim condition (ExIC) and the adequate valuation condition (AVC). That is, M is a *weak* epistemic game model in which the formula $\bigwedge_{\psi \in \Delta_0} \psi$ is true (i.e., $M, w \models \bigwedge_{\psi \in \Delta_0} \psi$ for all w in M) iff M is an epistemic game model. Therefore, the class **WEGM** in which the formula $\bigwedge_{\psi \in \Delta_0} \psi$ is true coincides with the class **EGM**. \square

The following Proposition 4.19 highlights that, thanks to the universal modality \square , the logical consequence problem relative to the class **WEGM** can be reduced to the validity problem relative to the class **WEGM**. The proof of the proposition is trivial, as we just need to apply the definitions of validity and logical consequence relative to **WEGM**.

Proposition 4.19. *For every EDGL-formula φ and for every finite set Δ of EDGL-formulas, $\Delta \models_{WEGM} \varphi$ iff $\models_{WEGM} \square \bigwedge_{\psi \in \Delta} \psi \rightarrow \varphi$.*

The following Lemma 4.20 provides an axiomatization result for EDGL relative to the class **WEGM**.

Lemma 4.20. *The set of validities of the logic EDGL relative to the class **WEGM** is completely axiomatized by the groups of axioms (1),(2),(3),(4) and (5) and by the rules of inference (7) and (8) in Figure 4.5.*

Proof. The proof is divided into three steps.

Step 1 The first step consists in providing an alternative semantics for EDGL relative to the class of *enriched* weak epistemic game models (**EWEGM**). An *enriched* weak epistemic game model is a tuple $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, \sim, v \rangle$ where $\langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ is a weak epistemic game model and \sim is an equivalence relation on W such that:

(C1) for all $i \in N$, $\sim_i \subseteq \sim$;

(C2) for all $u_1, \dots, u_n \in W$: if $u_i \sim u_j$ for all $i, j \in \{1, \dots, n\}$ then $\sim_1(u_1) \cap \dots \cap \sim_n(u_n) \neq \emptyset$.

The truth conditions of EDGL formulas relative to the class **EWEGM** are exactly like the truth conditions of EDGL formulas relative to the classes **WEGM** and **EGM**, except for the universal modality \Box that is interpreted as follows. Let $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, \sim, v \rangle$ be a **EWEGM** and let w be a world in M . Then:

- $M, w \models \Box\varphi$ iff $\forall w'$ such that $w \sim w'$, $M, w' \models \varphi$

Step 2 The second step consists in proving that the set of validities of EDGL relative to the class **EWEGM** is completely axiomatized by the groups of axioms (1),(2),(3),(4) and (5) and by the rules of inference (7) and (8) in Figure 4.5.

It is a routine task to check that all principles in Figure 4.5 except Axiom (5d) correspond one-to-one to their semantic counterparts on the models in the class **EWEGM**. This can be easily checked by using the existing algorithm SQEMA [Conradie et al., 2006]: for every axiom in Figure 4.5, it allows us to compute the corresponding first-order condition on the models in the class **EWEGM**.

In particular, the group of axioms (2) together with the inference rule (8) correspond to the fact that \sim is an equivalence relations.³ The group of axioms (3) corresponds to the fact that \sim_i is an equivalence relation, while Axiom (4b) corresponds to the seriality of the relation \rightarrow_i .

Remark 4.21. Note that the necessitation rules for the knowledge modality (i.e., from φ infer $K_i\varphi$) and for the belief modality (i.e., from φ infer $B_i\varphi$) do not need to be added to the axiomatization, as they are provable by Axioms (5a), Axiom (5b), the inference rule (7) and the inference rule (8).⁴

³Specifically, Axiom (2b) corresponds to reflexivity of the relation \sim , Axiom (2c) to transitivity and Axiom (2d) to symmetry.

⁴Suppose φ . Hence, by the inference rule (8), we infer $\Box\varphi$. Thus, by Axiom (5a) and the inference rule (7), we infer $K_i\varphi$. Furthermore, by Axiom (5b) and the inference rule (7), we infer $B_i\varphi$.

As to the group of axioms (5) we have the following correspondences: Axiom (5a) corresponds to the preceding condition C1: $\sim_i \subseteq \sim$; Axiom (5b) corresponds to the following condition in the definition of epistemic game model: $\rightarrow_i \subseteq \sim_i$; Axiom (5c) corresponds to the following condition: for all $w, v \in W$, if $w \sim_i v$ then $\rightarrow_i(v) \subseteq \rightarrow_i(w)$. Because of the reflexivity of the relation \sim_i the latter is equivalent to the following condition in the definition of epistemic game model: for all $w, v \in W$, if $w \sim_i v$ then $\rightarrow_i(v) = \rightarrow_i(w)$.

As to Axiom (5d) a bit more work is required. First of all, it is a routine task to verify that, in terms of correspondence theory, Axiom (5d) corresponds to the following condition:

(C2*) for all $w, u_1, \dots, u_n \in W$: if $u_1, \dots, u_n \in \sim(w)$ then there is $v \in W$ such that $v \in \sim(w)$ and $\sim_i(v) \subseteq \sim_i(u_i)$ for all $i \in N$.

Secondly, one can prove that the condition C2* and the condition C2 are equivalent. Let us prove first that C2 implies C2*. Suppose that $u_1, \dots, u_n \in \sim(w)$. Hence, by condition C2 and the fact that \sim is an equivalence relation, $\sim(w) \cap \sim_1(u_1) \cap \dots \cap \sim_n(u_n) \neq \emptyset$. It follows that there exists $v \in W$ such that $v \in \sim(w)$, $v \in \sim_1(u_1), \dots, v \in \sim_n(u_n)$. Since every \sim_i is an equivalence relation, for every $i \in N$, if $v \in \sim_i(u_i)$ then $\sim_i(v) = \sim_i(u_i)$. Thus, we can conclude that there exists $v \in W$ such that $v \in \sim(w)$ and $\sim_i(v) \subseteq \sim_i(u_i)$ for all $i \in N$.

Now let us prove that C2* implies C2. Suppose that $u_i \sim u_j$ for all $i, j \in \{1, \dots, n\}$. It follows that $u_1, \dots, u_n \in \sim(w)$ for some w . Hence, by condition C2*, there are $w, v \in W$ such that $v \in \sim(w)$ and $\sim_i(v) \subseteq \sim_i(u_i)$ for all $i \in N$. Since every \sim_i is an equivalence relation, for every $i \in N$, if $\sim_i(v) \subseteq \sim_i(u_i)$ then $\sim_i(v) = \sim_i(u_i)$. Thus, we can conclude that there are $w, v \in W$ such that $v \in \sim(w)$ and $\sim_i(v) = \sim_i(u_i)$ for all $i \in N$. By the fact that every relation \sim_i is reflexive, it follows that there are $w, v \in W$ such that $v \in \sim(w)$ and $v \in \sim_i(u_i)$ for all $i \in N$. Hence, $\sim_1(u_1) \cap \dots \cap \sim_n(u_n) \neq \emptyset$.

It is routine, too, to check that all principles given in Figure 4.5 are in the so-called Sahlqvist class [Sahlqvist, 1975]. Thus, because of the general Sahlqvist completeness theorem (cf. [Blackburn et al., 2001, Theorem 3.54]), they are complete with respect to the defined model classes.

Step 3 The third step consists in proving that the EDGL-semantics relative to the class **EWEGM** and the EDGL-semantics relative to the class **WEGM** are equivalent. Specifically, we show that for every EDGL-formula formula φ , φ is satisfiable in the class **WEGM** iff φ is satisfiable in the class **EWEGM**.

(\Rightarrow) Let us prove the left-to-right direction. Suppose φ is satisfiable in the class **WEGM**. This means that there is a *weak* epistemic game model $M = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, v \rangle$ and a world $w \in W$ such that $M, w \models \varphi$. We can build a corresponding *enriched* weak epistemic game model $M' = \langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, \sim, v \rangle$ with $\sim = W \times W$. Clearly, $M', w \models \varphi$.

(\Leftarrow) Let us prove the right-to-left direction. Suppose φ is satisfiable in the class **EWEGM**. This means that there is an *enriched* weak epistemic game model $M =$

$\langle W, \sim_1, \dots, \sim_n, \rightarrow_1, \dots, \rightarrow_n, \sim, v \rangle$ and a world $w \in W$ such that $M, w \models \varphi$. We can now build a *weak* epistemic game model $M' = \langle W', \sim'_1, \dots, \sim'_n, \rightarrow'_1, \dots, \rightarrow'_n, v' \rangle$ that corresponds to M . The construction of M' is made in two steps. We first consider the submodel $M_w = \langle W_w, \sim_{1,w}, \dots, \sim_{n,w}, \rightarrow_{1,w}, \dots, \rightarrow_{n,w}, \sim_w, v_w \rangle$ generated from M and w (cf. [Blackburn et al., 2001, Definition 2.5]): by the generated submodel property (cf. [Blackburn et al., 2001, Proposition 2.6]) we have $M_w, w \models \varphi$. M_w is also an *enriched* weak epistemic game model, and $\sim_w = W_w \times W_w$. The latter means that the operator \Box is interpreted as a universal modal operator. Finally, we can define $M' = \langle W', \sim'_1, \dots, \sim'_n, \rightarrow'_1, \dots, \rightarrow'_n, v' \rangle$ as follows:

- $W' = W_w$;
- for every $i \in N$, $\sim'_i = \sim_{i,w}$ and $\rightarrow'_i = \rightarrow_{i,w}$;
- $v' = v_w$.

It is a routine task to check that M' is indeed a *weak* epistemic game model and, by induction on the structure of φ , that we have $M', w \models \varphi$.

Lemma 4.20 is a consequence of: (i) the equivalence between the EDGL-semantics relative to the class **EWEGM** and the EDGL-semantics relative to the class **WEGM** (proved in Step 3) and, (ii) the completeness result for EDGL relative to the class **EWEGM** (proved in Step 2). \square

The last element we need for proving Theorem 4.17 is the following Proposition 4.22. Let $\vdash_{EDGL} \varphi$ and $\Vdash_{EDGL} \varphi$ mean, respectively, that the EDGL-formula φ is provable via the groups of axioms (1),(2),(3),(4), (5) and (6) and the rules of inference (7) and (8) in Figure 4.5 and that the EDGL-formula φ is provable via the groups of axioms (1),(2),(3),(4) and (5) and the rules of inference (7) and (8) in Figure 4.5.

Proposition 4.22. *For every EDGL-formula φ , if $\Vdash_{EDGL} \Box \bigwedge_{\psi \in \Delta_0} \psi \rightarrow \varphi$ then $\vdash_{EDGL} \varphi$, where Δ_0 is defined as in Proposition 4.18.*

Proof. Suppose $\Vdash_{EDGL} \Box \bigwedge_{\psi \in \Delta_0} \psi \rightarrow \varphi$. Hence, $\vdash_{EDGL} \Box \bigwedge_{\psi \in \Delta_0} \psi \rightarrow \varphi$.

By the inference rule (8) (viz. necessitation for \Box) and the group of axioms (6), we have $\vdash_{EDGL} \bigwedge_{\psi \in \Delta_0} \Box \psi$. By Axiom 2(a), we can derive $\vdash_{EDGL} \Box \bigwedge_{\psi \in \Delta_0} \psi$. Consequently, by the inference rule (7) (viz. modus ponens), we have that $\vdash_{EDGL} \varphi$. \square

Propositions 4.18, 4.19 and 4.22 together with Lemma 4.20 are sufficient to prove Theorem 4.17.

Suppose that $\models_{EGM} \varphi$. Hence, by Proposition 4.18 and Proposition 4.19, $\models_{WEGM} \Box \bigwedge_{\psi \in \Delta_0} \psi \rightarrow \varphi$. By Lemma 4.20, it follows that $\Vdash_{EDGL} \Box \bigwedge_{\psi \in \Delta_0} \psi \rightarrow \varphi$. Hence, by Proposition 4.22, $\vdash_{EDGL} \varphi$. \square

As to the logic DGL, namely the fragment of the logic EDGL presented in Section 4.2, we have the following axiomatization result. We do not prove it here, as the proof follows the general lines of the proof of Theorem 4.17.

Theorem 4.23. *The set of validities of the logic DGL relative to the class of doxastic game models (DGM) is completely axiomatized by the groups of axioms (1), (4) and (6) and the inference rule (7) in Figure 4.5, plus the following axioms and rules of inference for the belief modality B_i :*

- $B_i\varphi \rightarrow B_iB_i\varphi$
- $\widehat{B}_i\varphi \rightarrow B_i\widehat{B}_i\varphi$
- From φ infer $B_i\varphi$

Note that differently from EDGL, Axioms 4 and 5 as well as the necessitation rule for the belief modality B_i must be added to the axiomatics, as they are not derivable from the other principles.

4.5 Probabilistic Extension

In this section we want to show how to extend our analysis about the equivalence between Kripke-style semantics and type space semantics for epistemic modal languages to a modal language with probabilistic beliefs $\mathcal{L}_{EPGL}(G)$, interpreted over a specific class of Kripke models called *epistemic-probabilistic game models* (EPGM). The resulting logic EPGL (*epistemic-probabilistic game logic*) will then be an extension of the logic EDGL, that was itself introduced as an extension of DGL. Consequently, this section follows an approach that is incremental with respect to the logics that have been introduced so far.

4.5.1 Language

Let us first introduce a probabilistic language $\mathcal{L}_{EPGL}(G)$ for epistemic-probabilistic game logic EPGL, defined by the grammar:

$$\varphi ::= pl_i(a_i) \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_i\varphi \mid P_i^r\varphi \mid \Box\varphi$$

where $a_i \in A_i$, $i \in N$ and $r \in \mathbb{Q}$.

The language $\mathcal{L}_{EPGL}(G)$ is the language $\mathcal{L}_{EDGL}(G)$ of EDGL introduced in Section 4.3 with the difference of replacing the qualitative belief operator B_i with the probabilistic belief operator P_i^r , where P_i^r has to be read as “player i believes with probability at least r that φ is true”. It is easy to see that this language is incremental relative to $\mathcal{L}_{EDGL}(G)$: the belief operator B_i of $\mathcal{L}_{EDGL}(G)$ is expressible in $\mathcal{L}_{EPGL}(G)$ by the operator P_i^1 .

4.5.2 Semantics

4.5.2.1 Type spaces

The interpretation of the new formulas $P_i^r \varphi$ in a type space T is as one could expect, and it is given by the following clause:

- $T, \vec{t} \models P_i^r \varphi$ iff $\forall \vec{t}' \in \vec{T}$ if $t'_i = t_i$ $\sum_{\vec{t}'_{-i}: T, \vec{t}'_{-i} \models \varphi} \beta_i(t_i)(\vec{t}'_{-i}) \geq r$.

All the other clauses are the same as in Section 4.3. As mentioned before, when $r = 1$ we can show the equivalence with respect to the clause for B_i . Indeed,

$$T, \vec{t} \models B_i \varphi$$

iff

$$\forall \vec{t}' \in \vec{T} \text{ if } t'_i = t_i \text{ and } \beta_i(t_i)(\vec{t}'_{-i}) > 0 \text{ then } T, \vec{t}' \models \varphi$$

iff

$$\forall \vec{t}' \in \vec{T} \text{ if } t'_i = t_i \text{ then } \sum_{\vec{t}'_{-i}: T, \vec{t}'_{-i} \models \varphi} \beta_i(t_i)(\vec{t}'_{-i}) = 1$$

iff

$$T, \vec{t} \models P_i^1 \varphi.$$

4.5.2.2 Epistemic-probabilistic models

When we come to consider how to interpret the language $\mathcal{L}_{EPGL}(G)$ over Kripke models we firstly need to endow the structure with probabilistic belief relations for the agents. The resulting Kripke model will be an *epistemic-probabilistic model*, formally defined as follows.

Definition 4.24 (Epistemic-probabilistic model). An epistemic-probabilistic model is a tuple $M = \langle W, \sim_1, \dots, \sim_n, \phi_1, \dots, \phi_n, v \rangle$ where W , \sim_i and v are defined as in Definition 4.8, and ϕ_i is a function mapping every world w in W into a probability distribution $\phi_{i,w}$ over the worlds in $\sim_i(w)$ that satisfies the following conditions:

- $\sum_{w' \in \sim_i(w)} \phi_{i,w}(w') = 1$;
- if $w \sim_i w'$ then $\phi_{i,w} = \phi_{i,w'}$.

In order to interpret the language $\mathcal{L}_{EPGL}(G)$ on epistemic-probabilistic models we want to identify a subclass of them that we call *epistemic-probabilistic game models* EPGM.

Definition 4.25 (Epistemic-probabilistic game model). An epistemic-probabilistic game model EPGM is an epistemic-probabilistic model that satisfies conditions AVC and EIC in Definitions 4.2 and 4.9 plus the following condition ExIC':

- Ex interim condition revisited (ExIC’):

$$\forall i \in N, \forall w, w' \in W, \forall a_i \in A_i, \text{ if } w \sim_i w' \text{ and } pl_i(a_i) \in v(w) \\ \text{then } pl_i(a_i) \in v(w').$$

Now we can give the conditions to interpret the language $\mathcal{L}_{EPGL}(G)$ over epistemic-probabilistic game models EPGM. Apart from P_i^r formulas, the truth conditions for all the other formulas remain the same as for EDGM. The clause for P_i^r formulas is as follows:

- $M, w \models P_i^r \varphi$ iff $\sum_{w' \in \sim_i(w): M, w' \models \varphi} \phi_{i,w}(w') \geq r$.

4.5.3 Correspondence

It is also possible to establish correspondence results between type spaces and epistemic-probabilistic game models relative to the language $\mathcal{L}_{EPGL}(G)$. Given an arbitrary type space T , the corresponding epistemic-probabilistic model M^T is defined as the tuple $M^T = \langle W, (\sim_i)_{i \in N}, (\phi_{i,w})_{i \in N, w \in W}, v \rangle$ where:

- $W = \vec{T}$ is the set of worlds;
- $v : W \rightarrow 2^\Phi$ is the valuation function such that $pl_i(a_i) \in v(w)$ iff $\sigma_i(t_i) = a_i$;
- \sim_i is the accessibility relation of player i , given by: $\forall w, w' \in W, w \sim_i w'$ iff $t_i = t'_i$;
- $\phi_{i,w}$ is the probabilistic belief relation of player i at world w , defined as follows: $\phi_{i,w}(w') = r$ iff $\beta_i(t_i)(\vec{t}_{-i}) = r$ and $t_i = t'_i$.

It is easy to see that the Kripke model that we get after the transformation is an epistemic-probabilistic game model.

Proposition 4.26. *M^T is an epistemic-probabilistic game model.*

Proof. AVC and EIC conditions hold unchanged with respect to EGM. The only thing that differs from epistemic game models is ExIC’. Then, ExIC’ also holds in M^T since each type is associated with one and only one action and two worlds w and w' belong to the same partition cell of player i if and only if the type of i at w is the same as i ’s type at w' . \square

As for DGM and EGM before, we conclude the first half of the proof with a theorem showing that if a formula of $\mathcal{L}_{EPGL}(G)$ is satisfiable in a type space T , then it is also satisfiable in the corresponding EPGM M^T .

Theorem 4.27. *Let $\vec{t} \in \vec{T}$ be a state in T and let $w \in W$ be the corresponding world in M^T . For any φ in $\mathcal{L}_{EPGL}(G)$, if $T, \vec{t} \models \varphi$ then $M^T, w \models \varphi$.*

Proof. By induction on the structure of φ . We prove just the case $\varphi = P_i^r \psi$ since the others are similar to Section 4.3. Suppose $T, \vec{t} \models P_i^r \psi$. Therefore, we have that $\sum_{\vec{t}'_{-i}: T, t'_i \models \psi} \beta_i(t_i)(\vec{t}'_{-i}) \geq r$. By inductive hypothesis, $\sum_{w' \in \sim_i(w): M, w' \models \psi} \phi_{i,w}(w') \geq r$, hence $M^T, w \models P_i^r \psi$. \square

In the second half of the proof the other direction is carried out: we show that given any epistemic-probabilistic game model M it is possible to associate with it the corresponding type space T^M , and that M and T^M are semantically equivalent with respect to $\mathcal{L}_{EPGL}(G)$.

Apart from the probabilistic belief relation $\phi_{i,w}$, the construction of the type space T^M is the same as for the case of epistemic game models. What we need to define here is the belief function β_i for player i , given the probabilistic belief relation $\phi_{i,w}$ of player i in M . Then, for an arbitrary $M = \langle W, (\sim_i)_{i \in N}, (\phi_{i,w})_{i \in N, w \in W}, \nu \rangle$, we can associate with it the corresponding type space $T^M = \langle T_1, \dots, T_n, \beta_1, \dots, \beta_n, \sigma_1, \dots, \sigma_n \rangle$ defined as follows:

- a type t_i of player i for each i 's partition cell $\sim_i(w)$ in M : $T_i = \{\sim_i(w) : w \in W\}$;
- an action a_i of t_i specified by ν : $\sigma_i(t_i) = a_i$ iff $pl_i(a_i) \in \nu(w)$;
- a belief function β_i such that $\beta_i(t_i)(\vec{t}'_{-i}) = r$ iff $\phi_{i,w}(w') = r$.

Theorem 4.28. *Let $w \in W$ be a world of M and $\vec{t} \in \vec{T}$ the corresponding state in T^M . For any φ in $\mathcal{L}_{EPGL}(G)$, if $M, w \models \varphi$ then $T^M, \vec{t} \models \varphi$.*

Proof. By induction on the structure of φ .

We show only the case $\varphi = P_i^r \psi$ since the others are already covered in Section 4.3. Suppose $M, w \models P_i^r \psi$. Then, by definition $\sum_{w' \in \sim_i(w): M, w' \models \psi} \phi_{i,w}(w') \geq r$. By inductive hypothesis, $\sum_{\vec{t}'_{-i}: T, t'_i \models \psi} \beta_i(t_i)(\vec{t}'_{-i}) \geq r$, where $t_i = t'_i$. Hence $T^M, \vec{t} \models P_i^r \psi$. \square

The following corollary then states that the class of type spaces \mathbf{T} and the class of epistemic-probabilistic game models \mathbf{EPGM} are equivalent semantics with respect to the language $\mathcal{L}_{EPGL}(G)$.

Corollary 4.29. *A formula φ of $\mathcal{L}_{EPGL}(G)$ is satisfiable in \mathbf{T} iff it is satisfiable in \mathbf{EPGM} .*

4.6 Conclusion

This work aimed at explicitly showing the formal relations between the two main structures that are used in the literature on epistemic game theory: type spaces and Kripke models. We started with a language for belief and we proceeded by

extending the language taken into account. We noticed that with respect to belief the relation can be carried through in a straightforward way, using a rather standard KD45-concept of belief. When we tried later to establish a correspondence with respect to a language with both belief and knowledge the situation turned out to be more complicated. Interpreting in the most obvious way a S5-concept of knowledge on type spaces gave rise to some interesting properties (as shown in Section 4.3.5) that the usual S5-knowledge operator does not have. This means in a sense that Kripke models are less demanding towards the interpretation of knowledge and more suitable to express both knowledge and belief at the same time. The kind of knowledge representable in Kripke models is broader and more general than the interpretation of knowledge in type spaces, whereas the belief representation is basically the same in both structures.

An important thing to stress here is also that we axiomatized a *qualitative* logic for type spaces. Since Kripke models are normally qualitative structures, we decided to establish a correspondence from a qualitative point of view in the first place. We then showed what are the qualitative logics (for belief and for both belief and knowledge) and the classes of qualitative Kripke models (i.e., doxastic game models and epistemic game models respectively) corresponding to type spaces (Section 4.4). We presented at first the result for a qualitative language with at most two different operators, B_i and K_i . Using the same line of reasoning we could aim at proving the result for other epistemic attitudes, like strong belief or defeasible knowledge ([Baltag and Smets, 2008]). What we need to do is to enrich the semantics on the side of Kripke models by means of plausibility orderings or rankings over worlds instead of simple belief sets (see [Baltag and Smets, 2008] or [Spohn, 2012]). By a further enrichment of Kripke models towards more fine-grained semantics in the end we introduced probabilistic Kripke models and we extended the correspondence result to the epistemic-probabilistic language $\mathcal{L}_{EPGL}(G)$ (Section 4.5).

Part II

Evolutionary Analysis

As economic theorists, we organize our thoughts using what we call models. The word “model” sounds more scientific than “fable” or “fairy tale” although I do not see much difference between them.

(A. Rubinstein)

5.1 The Behavioral Gambit

Classic evolutionary game theory, as we have seen in the previous chapter, looks at a single, fixed fitness game and focuses on the evolution of behavior for that game alone. Although single-game models can be useful to study a particular type of interaction in isolation (e.g., [Sinervo and Lively, 1996] use the rock-paper-scissors game to model the mating system in the side-blotched lizard), assuming that the entire biological dynamics driving the evolution of a population can be modeled as a fixed and single game would be an oversimplified description of the reality. In the case of side-blotched lizards investigated by [Sinervo and Lively, 1996], for example, the mating behavior is certainly a primary component for the evolutionary success of an individual, but it is presumably not the only one. Lizards’ strategies when hunting for food or hiding from predators are equally important factors in determining the fitness of different individuals. If we want to get closer to more realistic evolutionary dynamics, we have to incorporate a variety of possible interactions into our models.

A related shortcoming of classic single-game models is that the phenotypes under selection are nothing more than simple behavioral traits, and each player just represents a single action of the fixed fitness game. In other words, pure actions are the only things that evolve within a single-game model. In contrast, some argue for studying the evolutionary competition of general behavior-generating

mechanisms. [Fawcett et al., 2013], in particular, explicitly point out the “behavioral gambit” of standard evolutionary models, and write

By focusing on expressed behavior and neglecting the underlying mechanism, behavioral ecologists unwittingly adopt the behavioral gambit, extending the phenotypic gambit beyond its accepted remit.

[...] Natural environments are so complex, dynamic, and unpredictable that natural selection cannot possibly furnish an animal with an appropriate, specific behavior pattern for every conceivable situation it might encounter. Instead, we should expect animals to have evolved a set of psychological mechanisms which enable them to perform well on average across a range of different circumstances.

This line of thoughts is clearly reminiscent of the work on ecological rationality encountered in Chapter 2. Sometimes [Fawcett et al., 2013] even adopt the same terminology used by Gigerenzer and colleagues when talking about *homo heuristicus*:

To understand why that behavior has evolved, we have to consider the adaptive value of the psychological mechanism which controls it, in the kinds of environments the animal would normally encounter.

Still, there seems to be a slightly different understanding of these mechanisms in the two approaches. While the heuristics of the adaptive toolbox are different behavioral rules unrelated to each other, the psychological mechanism sought by [Fawcett et al., 2013] looks like a more general and abstract principle that governs distinct expressed behaviors in a systematic way (we will expand on this issue later).

Either way, the study of behavior-generating mechanisms would hardly fit into classic evolutionary game theory, given the minimal conception of environment delivered by single-game models. The competition between general choice principles requires a richer and extended notion of environment, in terms of a multiplicity of possible interactions, or the concept of choice principle would otherwise collapse into that of simple behavior, and the two would no longer be distinguishable.

5.2 Choice Principles: A General Discussion

The essence of a choice principle is to associate different decision situations with action choices. Chapters 2 and 3 presented an assortment of solitary and interactive decision situations. Roughly speaking, a choice principle, or decision criterion, is then a method to select a specific action for any of those decision problems. Theoretically, a simple heuristic like those in *homo heuristicus*' adaptive toolbox can be a choice principle. For example, a simple rule dictating “always

choose the first option that comes to your mind” is a decision criterion, in that it specifies an act for any possible decision situation. On the contrary, a rule dictating “always choose act II ” is not guaranteed to be a valid choice principle, because it would specify a choice only for some of the games in Chapter 3, but it would give no behavioral prescription in most decision problems of Chapter 2 as well as in the Traveler’s dilemma of Chapter 3 for instance. But in order to relate choice principles more tightly to the decision-theoretic literature introduced in Chapter 2, we take a closer look inside choice principles, and express them as a function of two different things: a subjective *utility* and a subjective *belief*. This perspective will carry a twofold contribution. On the one hand, we can work with decision criteria that are more general and less case-specific than those of Gigerenzer’s homo heuristicus.¹ On the other hand, we can bring together the literature on ecological rationality and the classic works in decision theory. In other terms, this approach will allow the study of the ecological rationality of general decision criteria.

In general, the simplest form of decision problem an agent can be faced with is one in which DM merely has to select an action a in order to achieve a reward $\pi(a)$, where the outcome $\pi(a)$ does not depend on anything else than the choice of the agent (in particular, it does not depend on any possible state of the world or any action of other agents). In such a situation, we generally deem an action a as rational if it maximizes the utility of the agent $u(\pi(a))$. An action a' is thus inferior to a competing alternative a if and only if $u(\pi(a')) \leq u(\pi(a))$.

However, in many realistic scenarios of our interest, the outcome achieved by the agent depends on extrinsic factors that are not under her control (all the results in decision theory presented in Chapter 2 fall under this category). In such situations, an agent’s action a does not in itself yield a unique outcome, but rather selects a certain outcome function $b \mapsto \pi(a, b)$, where $b \in B$ is a generic variable whose value is not chosen by the agent (in solitary decision problems b is a state variable chosen by nature, in interactive decision problem b represents some actions chosen by other agents). Depending on what value b happens to take, a certain action a may then prove to be either better or worse than an alternative a' in terms of utility for DM. However, without additional assumptions, there is no straightforward way of comparing the two outcome functions $b \mapsto \pi(a, b)$ and $b \mapsto \pi(a', b)$, so that the choice of action is no longer a simple maximization

¹We just observed that a simple heuristic like those suggested by Gigerenzer and colleagues is not always general enough to be applied in all possible circumstances, from which stems the necessity of providing homo heuristicus with a variety of different heuristics. But admitting a multiplicity of possible heuristics in the adaptive toolbox might rise a second-order issue: in case of a decision problem where more than one heuristic is applicable, how should DM choose between different heuristics? Are there second-order heuristics to deal with these cases, or does DM just pick one at random? The same problem may also occur at higher levels. Our approach instead wants to equip DM with a general and single choice principle that can be used throughout. A choice principle must thus be flexible enough to be fit for all decision situations that DM may possibly encounter.

problem.

By and large, a method for imposing an ordering on such a space of outcome functions is what we call a choice principle. A *choice principle*, or *decision criterion*, is a rule that dictates how DM should act when faced with a choice whose consequences are uncertain. In general, such a decision will depend both on the objective structure of the decision problem, and on DM's subjective preference over different possible outcomes and subjective beliefs about the realization of these outcomes. We have seen that the objective features of the situation can be formally specified in terms of an outcome function

$$\pi : A \times B \rightarrow X$$

which maps an action $a \in A$ and a situation $b \in B$ to an outcome $\pi(a, b)$. It is thus natural to think of a as a control variable whose value is chosen by the agent in question, while b represents some external factors which influence the agent's outcome but are outside her control. A choice principle is a function that combines DM's subjective utility over X and beliefs over B in order to pick an action:

$$\text{Choice principle: Utility} \times \text{Beliefs} \rightarrow \text{Actions.}$$

From Chapter 1 and Chapter 2, we know that there are many possible ways to deal with the uncertainty about the relevant external factors b . Depending on the type of uncertainty the agent has about the value of b , this uncertainty may be best described in terms of a probability distribution, a set of probability distributions, a plausibility ordering, a set of values, or something else.² In turn, different kinds of uncertainty call for different notions of uncertainty resolution. For instance, when the agent can quantify the uncertainty about b in terms of a probability distribution P , the standard definition of best reply refers to the work by [Savage, 1954], and state that an action a^* is a best reply if it maximizes expected utility,

$$a^* \in \operatorname{argmax}_{a \in A} \mathbb{E}_P[u(\pi(a, b))].$$

This is still the standard model for decisions under uncertainty in many branches of economics. Specifically, most of the literature in game theory, as we have partially seen in Chapter 3, is still limited to the case of Bayesian agents. (A few recent attempts to extend the scope of game theory beyond the Bayesian paradigm can be found in [Battigalli et al., 2015], [Kajii and Ui, 2005], and [Liu, 2015]). However, Ellsberg's examples, among many others, show that we cannot presuppose that DM is able to quantify the uncertainty in a probabilistic way. Others (e.g., [Gilboa et al., 2012], [Gilboa et al., 2009], and [Gilboa, 2015]) even question the normativity of the Bayesian paradigm, as we have seen in Chapter 1.

²Here we almost exclusively consider (precise or imprecise) probabilistic models, but there is extensive literature also on decision theory without probabilities (e.g., [Tan and Pearl, 1994], [Baltag et al., 2009], [Brafman and Tennenholtz, 2000], [Lehmann, 1996]).

When DM's beliefs are not Bayesian, she has to appeal to other choice principles. If, for example, the uncertainty is just represented as a set of values of b , then a possible and famous decision criterion dictates to choose the action that maximizes the minimum. According to this criterion, an action a^* is a best reply if

$$a^* \in \operatorname{argmax}_{a \in A} \min_{b \in B} u(\pi(a, b)).$$

A third option, combining the first two, may be to express the agent's uncertainty as a set Γ of probabilities over B . A suitable decision criterion for this case would be the maximinimization of expected utility, already introduced in Chapter 2. An action a^* is then a best reply if

$$a^* \in \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, b))].$$

Notice that this third representation has some formal advantages, in that it can include the first two as a special case. Indeed, when $\Gamma = \{P\}$, the decision criterion reduces to maximization of expected utility; when $\Gamma = \Delta(B)$, then it corresponds to simple maxmin.

To sum up, while the subjective preferences over outcomes is specified by means of a vNM utility function $u : X \rightarrow \mathbb{R}$ (in line with what we have seen in Chapters 2 and 3, and with the literature on evolution of preferences, e.g., [Dekel et al., 2007], [Alger and Weibull, 2013]), a subjective belief can still be expressed in many different ways. A choice principle \hat{a} , however, must be able to resolve the uncertainty and to associate the agent's utility and beliefs with an action choice. Examples of choice principles (some of which we already encountered) are:

1. **Simple Maxmin:**

$$\hat{a}(u, B) = \operatorname{argmax}_{a \in A} \min_{b \in B} u(\pi(a, b))$$

2. **Simple Regret Minimization:**

$$\hat{a}(u, B) = \operatorname{argmax}_{a \in A} \min_{b \in B} \{u(\pi(a, b)) - \max_{a' \in A} u(\pi(a', b))\}$$

3. **Maxmin Expected Utility:**

$$\hat{a}(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, b))]$$

4. **Expected Regret Minimization:**

$$\hat{a}(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \{\mathbb{E}_P[u(\pi(a, b))] - \max_{a' \in A} \mathbb{E}_P[u(\pi(a', b))]\}$$

5. **Maximax Expected Utility:**

$$\hat{a}(u, \Gamma) = \operatorname{argmax}_{a \in A} \max_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, b))]$$

6. **Laplace rule** (for B finite):

$$\hat{a}(u, B) = \operatorname{argmax}_{a \in A} \sum_{b \in B} \frac{1}{|B|} u(\pi(a, b))$$

7. **Expected Utility Maximization:**

$$\hat{a}(u, P) = \operatorname{argmax}_{a \in A} \mathbb{E}_P[u(\pi(a, b))].$$

As observed earlier, each of these principles may require a different quantification of the uncertainty, but some uncertainty representations are more general and adaptable than others. For instance, we can encompass all the choice principles listed above by expressing uncertainty in terms of a set of probabilities Γ . Indeed, both maxmin and maximax expected utility are equivalent to expected utility maximization when the set Γ is a singleton, and in turn expected utility maximization reduces to Laplace rule when Γ is a single uniform probability over B . Simple regret minimization is expressible by means of expected regret minimization with $\Gamma = \Delta(B)$, just as simple maxmin is expressible in terms of maxmin expected utility with $\Gamma = \Delta(B)$, and expected regret minimization also boils down to expected utility maximization when Γ is a singleton.

A crucial part in determining DM's choices is played by the subjective representation of the decision situation, namely the manner of forming preferences and beliefs about a possibly uncertain world. Indeed, in order to prescribe an action, a choice principle needs to be given a specific utility and belief as input. We call the pair of a subjective utility and a subjective belief the *subjective representation* of the decision problem.

Once the subjective representation is fixed, the essence of a choice principle is the qualitative *order* over possible choices induced by the procedure for resolving uncertainty. Consider for example the case of maxmin and maximax expected utility, both acting on the same subjective representation (u, Γ) . The two choice principles will rank the available options according to two different criteria, inducing two different orders over possible actions. Given utility $u : X \rightarrow \mathbb{R}$ and belief Γ , an expected utility maximinimizer would prefer action a over a' if

$$\min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, b))] \geq \min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a', b))],$$

whereas an agent choosing according to maximax expected utility would prefer a over a' if

$$\max_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, b))] \geq \max_{P \in \Gamma} \mathbb{E}_P[u(\pi(a', b))].$$

Hence, a choice principle is the way DM resolves uncertainty and picks an action, and involves two levels of subjectivity. Firstly, DM has to decide which outcomes she would prefer and how likely she considers each outcome. This amounts to coming up with a subjective representation consisting of a subjective utility and a subjective belief. The second level fixes the method of resolving uncertainty and determines a ranking over possible acts.

5.3 On the Rationality of Choice Principles

Chapter 1 and Chapter 2 presented the basics of rational choice theory. We saw that the standard definition of rationality in textbooks and papers in economics and decision theory still traces back to the works by de Finetti [de Finetti, 1931], von Neumann and Morgenstern [von Neumann and Morgenstern, 1944], and Savage [Savage, 1954]. This tradition considers DM rational if she maximizes subjective expected utility (SEU). According to this point of view, DM is always able to quantify uncertainty in a probabilistic way, and then she maximizes her expected utility. Expected utility is subjective in the sense that it is a function of a subjective probabilistic belief P and a subjective utility u of the decision maker. We already noticed in Chapter 1 and Section 2.1 that SEU represents a weak and internal notion of rationality, whose unique requirement is DM's consistency of choices. To wit, a choice can be rational (i.e., the choice that maximizes subjective expected utility from DM's point of view), even if based on peculiar beliefs and/or aberrant preferences. On the other hand, it requires DM to always have probabilistic beliefs, and it dictates to always use expected utility maximization as decision criterion.

If beliefs and preferences are subjective, however, there is room for *rationalization* or *redescriptionism* of observable behavior. For example, in the case of interactive decision making, including considerations of fairness allows us to describe as rational (according to SEU) empirically observed behavior, such as in experimental prisoner's dilemmas or public goods games, that might otherwise appear irrational (e.g., [Fehr and Schmidt, 1999], [Charness and Rabin, 2002]).

The main objection to redescriptionism is that, without additional constraints, the notion of rationality is likely to collapse, as it seems possible to deem rational almost everything that is observed, given the freedom to adjust beliefs and preferences at will. An opposite position therefore emphasizes that there are many ways in which (i) the ascription of beliefs and preferences, and (ii) the uncertainty resolution method should be constrained by normative considerations of rationality as well. Subjective beliefs should be justified by evidence, and quantified in terms of a single probability function P only if possible; subjective preferences should be oriented towards tracking objective fitness. For instance, profit maximization seems a necessary requirement for evolution in a competitive market because only firms behaving according to profit maximization will survive in the

long run ([Alchian, 1950], [Friedman, 1953]). Once subjective beliefs and subjective utilities have been tuned as much as possible with objective chance and objective fitness respectively, the SEU paradigm prescribes that DM shall choose the action that maximizes expected utility.

An alternative view on rationality of choice is *adaptationism* ([Anderson, 1991], [Chater and Oaksford, 2000], [Hagen et al., 2012]). Adaptationism aims to explain rational behavior by appealing to evolutionary considerations: DMs have acquired choice principles that have proven to be adaptive with respect to the variable environment where they have evolved. The focus on the ecological rationality of different heuristics and choice mechanisms maintained by Gigerenzer’s school and the research program pursued by [Fawcett et al., 2013] both share this adaptive perspective on the study of rational choice. Our approach is along the same lines, in that we do not want to assess the rationality of choices on the basis of some system of axioms or aprioristic philosophical intuition about the nature of rationality. Our intuition is that the quality of a choice cannot be evaluated independently of the environment in which it takes place. In the following, we will specifically investigate the evolutionary fitness of different ways of forming subjective representations and resolving uncertainty across multiple decision problems.

5.4 The Game of Life

If agents deal with a rich and variable environment, they have to face many different choice situations. But, as noted earlier, standard evolutionary game models frequently simplify reality in at least two ways. Firstly, the environment is represented as a fixed fitness game; secondly, the focus of evolutionary selection is behavior for that stage game alone. This section instead introduces a general *multi-game* model that aims at overcoming the limitations of single-game models presented in previous sections, and conservatively extends the scope of evolutionary game theory to deal with evolutionary selection of general choice principles.³ The model will thus include a multiplicity of different fitness games,

³In particular, similar ideas recently appeared in [Zollman, 2008], [Bednar and Page, 2007], and [Robalino and Robson, 2016]. [Bednar and Page, 2007] use a multi-game framework, composed of a fixed selection of six possible games, to study the emergence of different cultural behaviors, and model agents as finite-state automata playing games from the fixed selection. [Zollman, 2008] explains seemingly “irrational” fair behavior in social dilemmas (like the Ultimatum game) by means of a model where agents have to play the Ultimatum game together with the Nash bargaining game, but they are constrained to choose the same strategy for both games. Finally, [Robalino and Robson, 2016] consider, in a more decision-theoretic setting, what subjective utility function a cognitively limited agent should be endowed with in order to maximize her evolutionary fitness. Our framework can then be viewed as a generalization of those models, mainly in that here players do not experience any cognitive constraint, and we allow for larger and possibly variable classes of games.

each representing a possible interaction between members of the population, and will capture the evolutionary dynamics determined by the multiplicity of interactions in the environment. For this reason, the multi-game model might also be called *the big game*, or *the game of life*. To illustrate the usefulness of the model, this chapter shows how it can enable us to investigate which choice principles are ecologically valuable and lead to high fitness, and which principles would instead be disfavored by natural selection in multi-game environments.

Research on the evolutionary selection of different subjective representations (especially in terms of different subjective utilities) has been subject of recent interest in theoretical economics, giving rise to a body of literature under the name of *evolution of preferences* (see [Samuelson, 2001], [Robson and Samuelson, 2011]). There, the phenotypes under selection are usually preferences over outcomes, modeled as different subjective utility functions $u : X \rightarrow \mathbb{R}$. Our framework generalizes those models, and, in addition to the selection of subjective preferences, it also allows to study the evolution of different subjective beliefs and decision criteria. In fact, a crucial point that we are going to show is that questions of preference evolution should take variability in uncertainty representation and uncertainty resolution into account as well. We demonstrate that, particularly when agents have *imprecise* probabilistic beliefs (see [Dempster, 1967], [Levi, 1974], [Shafer, 1976], [Gardenfors and Sahlin, 1982]), the way they resolve uncertainty to an action choice is a fundamental issue for determining evolutionary selection. The result is relevant in that it also offers an evolutionary comparison in terms of ecological rationality between the main decision criteria presented in the chapter on rational choice.

5.5 The Model

We denote by \mathcal{G} the set of fitness games that can possibly be played in a given population. For simplicity, we assume that all fitness games $G = \langle N, X, (A_i, \Phi_i)_{i \in N}, \pi \rangle$ are symmetric two-player games, i.e., such that $N = \{1, 2\}$, $A_1 = A_2$ and $\Phi_1(\pi(a_1, a'_2)) = \Phi_2(\pi(a'_1, a_2)) =: \Phi(\pi(a, a'))$.⁴ As argued in Section 3.3, it is reasonable to assume that the players' uncertainty about the distribution of different phenotypes in the population may be non-probabilistic. To allow for this case, we represent uncertainty by means of a set Γ of probability distributions. Given the game-theoretic setting, the subjective belief Γ is now a set of probability functions over the co-player's actions, $\Gamma \subseteq \Delta(A)$.

Overall, in the present context a phenotype is a triple (\hat{a}, u, Γ) consisting of a decision criterion $\hat{a} : u \times \Gamma \rightarrow A$, a subjective utility $u : X \rightarrow \mathbb{R}$, and a subjective belief $\Gamma \subseteq \Delta(A)$. For brevity let us denote phenotype $(\hat{a}^i, u^i, \Gamma^i)$ simply by t^i ,

⁴Since the game is symmetric, we can simply write $\Phi(\pi(a, a'))$ for $\Phi_1(\pi(a_1, a'_2))$ and $A := A_1 = A_2$, as in Chapter 3. Furthermore, notice that all definitions can be extended to more general cases.

and the set of all phenotypes in the population by \mathcal{T} . The fitness of a phenotype is then measured in terms of expected fitness. Formally, the fitness of phenotype $t^i = (\hat{a}^i, u^i, \Gamma^i)$ against phenotype $t^j = (\hat{a}^j, u^j, \Gamma^j)$ in a symmetric two-player game $G = \langle \{1, 2\}, X, A, \Phi, \pi \rangle$ is given by:⁵

$$F_G(t^i, t^j) = \Phi(\pi(\hat{a}^i(u^i, \Gamma^i), \hat{a}^j(u^j, \Gamma^j))).$$

We want, however, to study the fitness and the evolutionary selection of different choice principles in a rich multi-game environment. Towards this end, fix a class \mathcal{G} of symmetric two-player fitness games, together with a probability measure P_G that specifies the occurrence probability of games $G \in \mathcal{G}$. Intuitively, the probability P_G encodes the statistical properties of the environment. The *game of life* is then a tuple $GoL = \langle \mathcal{T}, \mathcal{G}, P_G, F \rangle$, where \mathcal{T} is the set of phenotypes, \mathcal{G} is the set of possible games, P_G is the probability distribution of possible games, and $F : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ is the (meta-)fitness function, defined as:

$$F(t^i, t^j) = \int F_G(t^i, t^j) dP_G(G). \quad (5.1)$$

Hence, $F(t^i, t^j)$ determines the evolutionary fitness of phenotype t^i against phenotype t^j in the multi-game. It is also possible to compute the average fitness of phenotype t^i against the population, that is given by:

$$F(t^i) = \int F(t^i, t^j) dP_t(t^j) = \int \int F_G(t^i, t^j) dP_t(t^j) dP_G(G) \quad (5.2)$$

where $P_t(t^j)$ is the probability of encountering a co-player of phenotype t^j .

Multi-games are thus abstract models for the evolutionary competition between different choice principles and subjective representations in interactive decision making contexts. Standard notions of evolutionary game theory apply to multi-games as well. For example, a profile (t^i, t^i) is a strict Nash equilibrium of GoL if $F(t^i, t^i) > F(t^j, t^i)$ for all t^j ; a phenotype t^i is evolutionarily stable if for all t^j : (i) $F(t^i, t^i) \geq F(t^j, t^i)$ and (ii) $F(t^i, t^i) = F(t^j, t^i) \Rightarrow F(t^i, t^j) > F(t^j, t^j)$; it is neutrally stable if for all t^j : (i) $F(t^i, t^i) \geq F(t^j, t^i)$ and (ii) $F(t^i, t^i) = F(t^j, t^i) \Rightarrow F(t^i, t^j) \geq F(t^j, t^j)$. Similarly, evolutionary dynamics can be applied to multi-games. Later we will also turn towards a dynamical analysis in terms of replicator dynamics and replicator-mutator dynamics.

5.6 A Bold Alternative to the Alternatives

The model introduced in the previous section is rather general and abstract. In this section, we concretely compare a selection of subjective representations and

⁵Whenever a choice mechanism would not select a unique action, we assume that the player chooses one of the equally optimal actions at random. I.e., $F_G(t^i, t^j) = \sum_{a \in \hat{a}^i(u^i, \Gamma^i)} \sum_{a' \in \hat{a}^j(u^j, \Gamma^j)} \frac{1}{|\hat{a}^i(u^i, \Gamma^i)|} \frac{1}{|\hat{a}^j(u^j, \Gamma^j)|} \Phi(\pi(a, a'))$.

decision criteria against each other. As for subjective utilities, consider initially the *objective* utility u that coincides with the fitness $u(\pi(a, a')) = \Phi(\pi(a, a'))$ for all $G \in \mathcal{G}$ and all $a, a' \in A$. The objective utility perfectly tracks the evolutionary fitness of the outcomes, and looks like the utility that an agent should be willing to adopt from a normative point of view.⁶

For a start, the subjective beliefs that we take into consideration are two:

1. \tilde{P} , a *precise* uniform belief over the co-player's actions: $\forall a \in A, \tilde{P}(a) = \frac{1}{|A|}$;
2. $\tilde{\Gamma}$, a maximally *imprecise* belief over the co-player's actions: $\tilde{\Gamma} = \Delta(A)$.

We have seen that these two kinds of belief underlie two different views on uncertainty. Faced with uncertain events, a Bayesian agent will always form a precise belief, specified by a *single* probability P . In the absence of any information about future uncertain events, the Bayesian would most likely invoke the *principle of insufficient reason*, and choose a uniform probability over the states. In contrast, others have argued against the obligation of representing a belief by means of a single probability measure, opposite to the Bayesian paradigm (e.g., [Gilboa et al., 2012], [Gilboa and Marinacci, 2013]). They argue instead in favor of a more encompassing account, according to which uncertainty can be *unmeasurable*, i.e., not represented by a unique measure. This line of thought appears extremely relevant in game-theoretic contexts too. Indeed, in a recent paper, [Battigalli et al., 2015] write:

Such [unmeasurable] uncertainty is inherent in situations of strategic interaction. This is quite obvious when such situations have been faced only a few times. (p. 646)

In evolutionary game theory, players obviously face uncertainty about the composition of the large population that they are part of, and consequently about the (type of) co-player that they are randomly paired with at each round and about the co-player's choice. In case of complete lack of information about the composition of the population, a non-Bayesian player would thus entertain maximal unmeasurable uncertainty, i.e., a maximally imprecise belief.⁷ As already anticipated, we will see that the way agents form beliefs, and the possibility of holding imprecise beliefs in particular, can have a fundamental impact on their evolutionary success.

Finally, the decision criteria that we start with are two:

⁶But see also [Alger and Weibull, 2013], among others.

⁷Such a radical uncertainty could ensue, in other cases, if agents have no conception of their co-player or her preferences. Unsophisticated agents, as considered in evolutionary game theory, might be entirely unaware of the interactive decision situation that they are engaged in (see [Heifetz et al., 2013], for game-theoretic models of unawareness). It is therefore not ludicrous to consider maximal uncertainty first and tend to more sophisticated ways of forming beliefs later (more on this below).

1. **Maxmin Expected Utility:**

$$Mm(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, a'))];$$

2. **Expected Regret Minimization:**

$$Reg(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \{ \mathbb{E}_P[u(\pi(a, a'))] - \max_{a'' \in A} \mathbb{E}_P[u(\pi(a'', a'))] \}.$$

As motivation for this comparison, we know from Chapter 2 that regret minimization is one of the main alternatives to utility maximization in decision theory, to the point that has been recently called “a bold alternative to the alternatives” ([Bleichrodt and Wakker, 2015]). The notion of regret in decision theory dates back at least to the work by [Savage, 1951], and has later been developed by [Loomes and Sugden, 1982, Fishburn, 1982, Bell, 1982] independently. Recently, [Halpern and Pass, 2012] showed how the use of regret minimization can give solutions to game-theoretic puzzles (like the Traveller’s dilemma and the Centipede game) in a way that is closer to everyday intuition and empirical data (see [Rubinstein, 2006]).

Some facts follow from the selection of decision criteria and subjective representations under consideration. The first is related to our focus on different types of uncertainty that players may entertain.

Fact 5.1. *For any precise (Bayesian) belief $\Gamma = \{P\}$ and subjective utility u , maxi(mini)mization of expected utility and minimization of expected regret are behaviorally equivalent. I.e., when Γ is a singleton set, $Mm(u, \Gamma) = Reg(u, \Gamma)$ for any subjective utility u .*

The following observations highlight that, by combining the selection of subjective beliefs and choice principles considered here, it is possible to embrace other choice principles listed previously. The next fact expresses a behavioral equivalence which we will make use of in the following section.

Fact 5.2. *In the class of 2×2 symmetric games (i.e., symmetric two-player games with two acts for each player), the acts selected by Laplace rule are exactly the acts selected by simple regret minimization.*

Fact 5.3 states two behavioral equivalences in the case of maximally imprecise beliefs.

Fact 5.3. *Fix a subjective utility u . The acts selected by simple maxmin are exactly the acts selected by maxmin expected utility for maximally imprecise beliefs $\tilde{\Gamma}$. Analogously, the acts selected by simple regret minimization are exactly the acts selected by expected regret minimization for maximally imprecise beliefs $\tilde{\Gamma}$.*

In light of Fact 5.3, $Mm(\Phi, \tilde{\Gamma})$ coincides with simple maximinimization of objective fitness, and likewise $Reg(\Phi, \tilde{\Gamma})$ is equivalent to simple regret minimization with respect to objective fitness.

A simple example may be helpful to show how choice principles and subjective representations work in practice. Consider the fitness game of Section 3.3 again. The next table pictures Ann's utility when it coincides with objective fitness Φ .

	I	II
I	2	1
II	0	5

While simple maxmin will choose action I , simple regret minimization will select II . Indeed, action I has the higher minimum, $1 > 0$, but it also has lower negative regret, $1 - 5 < 0 - 2$, and, hence, higher regret.

5.7 Results

5.7.1 Simulation results

Since for now we keep all subjective utilities u fixed to the evolutionary fitness Φ , different action choices of different players will only depend on differences in the decision criterion \hat{a} and/or subjective beliefs Γ . In other words, two phenotypes $(\hat{a}^i, u^i, \Gamma^i)$ and $(\hat{a}^j, u^j, \Gamma^j)$ can only differ in two of the three components, since we fix $u^i = u^j = \Phi$. Consequently, a phenotype $t^i \in \mathcal{T}$ is now fully specified by the pair (\hat{a}^i, Γ^i) , and we will directly refer to pairs like $(Reg, \tilde{\Gamma})$ or (Mm, \tilde{P}) as the *type* of the player for brevity. Sometimes we will also distinguish types by referring only to the choice principle or the subjective belief: for instance, $(Reg, \tilde{\Gamma})$ and (Reg, \tilde{P}) are regret types, while $(Mm, \tilde{\Gamma})$ and $(Reg, \tilde{\Gamma})$ are imprecise types.

As observed earlier, multi-games factor in statistical properties of the environment. For particular empirical purposes, one could consult a specific class of games \mathcal{G} with appropriate, maybe empirically informed probability P_G in order to match the natural environment of a given population. For our present purposes, let \mathcal{G} be a set of symmetric two-player fitness games with two acts for a start. Each game $G \in \mathcal{G}$ is then individuated by a quadruple of numbers $G = (a, b, c, d)$, as shown in the following table.

	I	II
I	$a; a$	$b; c$
II	$c; b$	$d; d$

As for the occurrence probability $P_G(G)$ of game G , we imagine that the values a, b, c, d are i.i.d. random variables sampled from the set $\{0, \dots, 10\}$ according to uniform probability P_V . Using Monte Carlo simulations, we can then approximate

	$(Reg, \tilde{\Gamma})$	$(Mm, \tilde{\Gamma})$	(Reg, \tilde{P})	(Mm, \tilde{P})
$(Reg, \tilde{\Gamma})$	6.663	6.662	6.663	6.663
$(Mm, \tilde{\Gamma})$	6.486	6.484	6.486	6.486
(Reg, \tilde{P})	6.663	6.662	6.663	6.663
(Mm, \tilde{P})	6.663	6.662	6.663	6.663

Table 5.1: Average evolutionary fitness from Monte Carlo simulations of 100.000 symmetric 2×2 games.

the values of equation (5.1) to construct the meta-fitness for the game of life. Results based on 100.000 randomly sampled games are given in table 5.1.⁸

Simulation results obviously reflect Fact 5.1 and Fact 5.2 in that all encounters in which types $(Reg, \tilde{\Gamma})$, (Reg, \tilde{P}) or (Mm, \tilde{P}) are substituted for one another yield identical results. More interestingly, Table 5.1 shows that $(Mm, \tilde{\Gamma})$, the simple maxmin strategy, is strictly dominated by the three other types: in each column (i.e., for each type of co-player), simple maxmin is strictly worse than any of the three competitors. This has a number of interesting consequences.

If we restrict attention to subjective representations with imprecise beliefs only, then a monomorphic state in which every agent has regret-based choice principle is the only *evolutionarily stable state*. More strongly, since $(Mm, \tilde{\Gamma})$ is strictly dominated by $(Reg, \tilde{\Gamma})$, we expect selection that is driven by (expected) fitness to invariably weed out maxmin players $(Mm, \tilde{\Gamma})$ in favor of regret minimizers $(Reg, \tilde{\Gamma})$. In terms of decision criteria, this means that simple regret minimization is evolutionarily better than simple maxmin over the class of games considered.

Next, if we look at the competition between all four types represented in Table 5.1, $(Reg, \tilde{\Gamma})$ is no longer evolutionarily stable. Given the behavioral equivalences of Fact 5.1 and Fact 5.2, types $(Reg, \tilde{\Gamma})$, (Reg, \tilde{P}) , and (Mm, \tilde{P}) are all *neutrally stable*. But since $(Mm, \tilde{\Gamma})$ is strictly dominated and so disfavored by fitness-based selection, we are still drawn to conclude that simple maxmin behavior is weeded out in favor of a population with a random distribution of the remaining three types.

Simulation results of the (discrete time) *replicator dynamics* indeed show that random initial population configurations are attracted to states with only three player types: $(Reg, \tilde{\Gamma})$, (Reg, \tilde{P}) and (Mm, \tilde{P}) . The relative proportions of these depend on the initial shares in the population. This variability fully disappears

⁸Concretely, 100.000 games were sampled repeatedly by choosing independently four integers between 0 and 10 uniformly at random. For each game, the action choices of all four types were determined and fitness from all pairwise encounters recorded. The number in each cell of Table 5.1 is the average fitness for the type listed in the row when matched with the type in the column.

if we add a small mutation rate to the dynamics. Take a fixed, small mutation rate ϵ for the probability that a player's decision criterion *or* her subjective belief changes to another criterion or belief. The probability that a player's type randomly mutates into a completely different type with altogether different decision criterion and different belief would then be ϵ^2 . With these assumptions about "component-wise mutations", numerical simulations of the (discrete time) *replicator-mutator dynamics* show that already for very small mutation rates almost all initial population states converge to a single fixpoint in which the majority of players have regret-based decision criterion. For instance, with $\epsilon = 0.001$, almost all initial populations are attracted to a final distribution with proportions:

$$\frac{\begin{array}{cccc} (Reg, \tilde{\Gamma}) & (Mm, \tilde{\Gamma}) & (Reg, \tilde{P}) & (Mm, \tilde{P}) \\ 0.289 & 0.021 & 0.398 & 0.289 \end{array}}{}$$

What this suggests is that, if biological evolution selects behavior-generating mechanisms, not behavior as such, it need not be the case that behaviorally equivalent mechanisms are treated equally all the while. If mutation probabilities are a function of individual components, it can be the case that certain components of such behavior-generating mechanisms are more strongly favored by a process of random mutation and selection. This is exactly the case of regret-based decision criterion. Since expected regret minimization is much better in connection with imprecise beliefs than maxmin expected utility is, the proportion of regret-based decision makers, and particularly of precise expected regret minimizers (Reg, \tilde{P}) , in the attracting state is substantially higher than that of expected utility maximizers, (Mm, \tilde{P}) , even though these types are behaviorally equivalent.

5.7.2 Analytical results

Results based on the single multi-game in Table 5.1 are not fully general and possibly spoiled by random fluctuations in the sampling procedure. Fortunately, for the case of 2×2 symmetric games, the main result that maxmin types $(Mm, \tilde{\Gamma})$ are strictly dominated by regret minimizers $(Reg, \tilde{\Gamma})$ can also be shown analytically for considerably general conditions.

Proposition 5.4. *Let \mathcal{G} be the class of 2×2 symmetric games $G = (a, b, c, d)$ generated by i.i.d. sampling a, b, c, d from a set of values with at least three elements in the support. Then, $(Reg, \tilde{\Gamma})$ strictly dominates $(Mm, \tilde{\Gamma})$ in the resulting meta-game.*

Proof. All proofs are in Appendix 5.10. □

Corollary 5.5. *Let \mathcal{G} be as in Proposition 5.4. If we only consider imprecise types, $(Mm, \tilde{\Gamma})$ and $(Reg, \tilde{\Gamma})$, then the unique evolutionarily stable state of RD is a monomorphic population of type $(Reg, \tilde{\Gamma})$.*

This shows that the main conclusions drawn in the previous section based on the approximated multi-game of Table 5.1 hold more generally for arbitrary 2×2 symmetric games with i.i.d. sampled fitness.

5.8 Extensions

How do the basic results from the preceding section carry over to richer models? Section 5.8.1 first introduces further conceptually interesting phenotypes that have been considered in the evolutionary game theory literature. Section 5.8.2 then addresses the case of symmetric two-player $n \times n$ games for $n \geq 2$. Finally, Section 5.8.3 presents a brief comparison to the case of solitary decision making.

5.8.1 More preference types

Given a game G , any function $u : X \rightarrow \mathbb{R}$ is a possible subjective preference. We have so far considered only objective preferences (i.e., $u^i = u^j = \Phi$ for all $t^i, t^j \in \mathcal{T}$ and all $G \in \mathcal{G}$), but the space of possible subjective preferences is enormous. Let us now look at other relevant subjective preferences that have been investigated, especially in behavioral economics and in evolutionary game theory. A famous example is the *altruistic* preference (e.g., [Becker, 1976], [Bester and Güth, 1998]), summoned to explain the possibility of altruistic behavior. At the other end of the spectrum, the *competitive* preference is located. The corresponding utilities are defined as follows:

1. *altruistic* utility:⁹ for all $G \in \mathcal{G}$,

$$\text{alt}(\pi(a, a')) := \Phi(\pi(a, a')) + \Phi(\pi(a', a));$$

2. *competitive* utility: for all $G \in \mathcal{G}$,

$$\text{com}(\pi(a, a')) = \Phi(\pi(a, a')) - \Phi(\pi(a', a)).$$

Once different subjective utilities have been introduced, there is variability in the players' subjective preferences too. Consequently, a player's type t^i can no longer be identified by the pair (\hat{a}^i, Γ^i) only, but needs the complete triple $(\hat{a}^i, u^i, \Gamma^i)$ to be fully specified. Throughout this section, a type t^i will thus correspond to a triple $(\hat{a}^i, u^i, \Gamma^i)$.

Table 5.2 shows results of Monte Carlo simulations that approximate the expected fitness in the game of life with all the subjective preferences considered

⁹A more general formulation would be to define γ -altruistic utility, for $\gamma \in [0, 1]$, $\text{alt}_\gamma(a, a') := \Phi(\pi(a, a')) + \gamma\Phi(\pi(a', a))$. Since we are not interested in the evolution of degrees of altruism, here we simply fix $\gamma = 1$. Analogously for γ -competitive utilities too.

	$(Reg, \Phi, \tilde{\Gamma})$	$(Mm, \Phi, \tilde{\Gamma})$	$(Mm, com, \tilde{\Gamma})$	$(Mm, alt, \tilde{\Gamma})$	(Reg, Φ, \tilde{P})	(Mm, Φ, \tilde{P})	(Mm, com, \tilde{P})	(Mm, alt, \tilde{P})
$(Reg, \Phi, \tilde{\Gamma})$	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
$(Mm, \Phi, \tilde{\Gamma})$	6.486	6.484	6.088	6.703	6.486	6.486	6.088	6.875
$(Mm, com, \tilde{\Gamma})$	6.323	6.758	5.496	6.977	6.323	6.323	5.496	7.149
$(Mm, alt, \tilde{\Gamma})$	5.949	5.722	5.326	6.396	5.949	5.949	5.326	6.568
(Reg, Φ, \tilde{P})	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
(Mm, Φ, \tilde{P})	6.663	6.662	5.829	7.105	6.663	6.663	5.829	7.489
(Mm, com, \tilde{P})	6.323	6.758	5.496	6.977	6.323	6.323	5.496	7.149
(Mm, alt, \tilde{P})	6.331	5.893	5.497	6.566	6.331	6.331	5.497	7.152

Table 5.2: Average evolutionary fitness from Monte Carlo simulations of 100.000 symmetric 2×2 games with four preference types.

so far.¹⁰ These results confirm basic intuitions about altruistic and competitive types: everybody would like to have an altruistic co-player and nobody would like to play with a competitive player. Perhaps more surprisingly, $(Mm, alt, \tilde{\Gamma})$ comes up strictly dominated by $(Mm, com, \tilde{\Gamma})$, but competitive types themselves are worse off against all types except against maxmin players $(Mm, \Phi, \tilde{\Gamma})$ than any of the behaviorally equivalent types $(Reg, \Phi, \tilde{\Gamma})$, (Reg, Φ, \tilde{P}) , and (Mm, Φ, \tilde{P}) . It thus follows that the previous results still obtain for the larger game of life in Table 5.2: $(Reg, \Phi, \tilde{\Gamma})$, (Reg, Φ, \tilde{P}) , and (Mm, Φ, \tilde{P}) are still neutrally stable; simulation runs of the (discrete-time) replicator dynamics on the 8×8 multi-game from Table 5.2 end up in population states consisting of only these three types in variable proportions.

In sum, the presence of other subjective representations, such as those based on altruistic or competitive utilities, does not undermine, but rather strengthens our previous results.

5.8.2 More actions

Results from Section 5.7 relied heavily on Fact 5.2, which is no longer true when we look at arbitrary $n \times n$ games. Table 5.3 gives approximations of expected fitness in the class of $n \times n$ symmetric games. Concretely, the numbers in table 5.3 are averages of evolutionary fitness obtained in 100.000 randomly sampled symmetric games, where each game G was sampled by first picking a number of acts $|A| \in \{2, \dots, 10\}$ uniformly at random, and then filling the necessary $|A| \times |A|$ fitness matrix with i.i.d. sampled numbers, as before.

The most important result is that the regret minimizing type $(Reg, \Phi, \tilde{\Gamma})$ is strictly dominated by (Reg, Φ, \tilde{P}) and by (Mm, Φ, \tilde{P}) in the multi-game from Table 5.3. This means that while simple regret minimization can thrive in some evolutionary contexts, there are also contexts where it is demonstrably worse off. While this may be bad news for simple regret minimizers $(Reg, \Phi, \tilde{\Gamma})$, it is

¹⁰Notice that in the simulations altruistic and competitive types may have precise as well as imprecise beliefs, but, for reasons of space, here we paired altruistic and competitive preferences only with maxmin expected utility as decision criterion.

	$(Reg, \Phi, \tilde{\Gamma})$	$(Mm, \Phi, \tilde{\Gamma})$	$(Mm, com, \tilde{\Gamma})$	$(Mm, alt, \tilde{\Gamma})$	(Reg, Φ, \tilde{P})	$(Mm^\circ, \Phi, \tilde{P})$	(Mm, com, \tilde{P})	(Mm, alt, \tilde{P})
$(Reg, \Phi, \tilde{\Gamma})$	6.567	6.570	5.650	6.992	6.564	6.564	5.593	7.409
$(Mm, \Phi, \tilde{\Gamma})$	6.476	6.483	5.896	6.818	6.484	6.484	5.850	7.124
$(Mm, com, \tilde{\Gamma})$	6.468	6.647	5.512	7.169	6.578	6.578	5.577	7.354
$(Mm, alt, \tilde{\Gamma})$	5.968	5.923	5.363	6.685	5.975	5.975	5.086	6.973
(Reg, Φ, \tilde{P})	6.908	6.918	5.988	7.456	6.929	6.929	5.934	7.783
(Mm, Φ, \tilde{P})	6.908	6.918	5.988	7.456	6.929	6.929	5.934	7.783
(Mm, com, \tilde{P})	6.529	6.680	5.445	7.276	6.542	6.542	5.521	7.440
(Mm, alt, \tilde{P})	6.450	6.337	5.772	6.978	6.457	6.457	5.479	7.500

Table 5.3: Average evolutionary fitness from Monte Carlo simulations of 100.000 symmetric $n \times n$ games with four preference types.

	(Reg, Φ)	(Mm, Φ)	(Mm, com)	(Mm, alt)
(Reg, Φ)	6.926	6.926	5.942	7.757
(Mm, Φ)	6.924	6.924	5.948	7.751
(Mm, com)	6.566	6.570	5.481	7.434
(Mm, alt)	6.463	6.461	5.478	7.469

Table 5.4: Meta-game for the evolutionary competition when beliefs are exogenously given (see main text).

not the case that regret types *as such* are weeded out by selection. Since, by Fact 5.1, (Reg, Φ, \tilde{P}) and (Mm, Φ, \tilde{P}) are behaviorally equivalent in general, it remains that selection in the game of life constructed from $n \times n$ games will still not eradicate regret-based principles.

On the other hand, there are plenty of ways in which the basic insights from Proposition 5.4 can make for situations in which evolution would favor regret types, even in $n \times n$ games. If, for example, the belief of a player is a trait that biological evolution has no bite on, but rather something that the particular choice situation would exogenously give us (possibly because of the different amount and quality of information available in different contexts), then expected regret minimizers can again drive out expected utility maximizers altogether. For example, suppose that only choice principles and preference representations compete, and that agents' beliefs are exogenously given in such a way that both players hold a precise (Bayesian) uniform belief with probability p and they both have a maximally imprecise belief otherwise. A phenotype is then a pair (\hat{a}, u) of a choice principle \hat{a} and a subjective utility u , while the beliefs are assigned game-by-game according to probability p . This transforms the multi-game from Table 5.3 into a simpler 4×4 meta-game in which the fitness obtained by type (\hat{a}, u) is the weighted average over the evolutionary fitness of the types including choice principle \hat{a} and preference u in Table 5.3. Setting $p = 0.98$ for illustration, we get the multi-game in Table 5.4.

The only evolutionarily stable state of this multi-game is again a monomorphic

population of regret types. Accordingly, all our simulation runs of the (discrete-time) replicator dynamics converge to monomorphic regret-type populations. The reason why expected regret minimizers prosper is that they have a substantial fitness advantage when paired with imprecise beliefs (Proposition 5.4). If unmeasurable uncertainty is exogenously given as something that happens to agents because of the information available in some choice situations, and even if that happens only very infrequently (i.e., for rather low p), regret types will outperform expected utility maximizers.

5.8.3 Solitary decisions

To see how different choice principles behave in evolutionary competition based on solitary decision making, we approximated, much in the spirit of multi-games, average accumulated fitness obtained in randomly generated solitary decision problems. For our purposes, a decision problem can be specified by a tuple $D = \langle S, X, A, \Phi, \pi \rangle$, where S is the finite set of states of the world, X is the set of outcomes, A is the finite set of acts, $\pi : S \times A \rightarrow X$ is the outcome function, and $\Phi : X \rightarrow \mathbb{R}$ is the fitness function.¹¹ We generate arbitrary decision problems by selecting, uniformly at random, numbers of states and acts $|S|, |A| \in \{2, \dots, 10\}$ and then filling the fitness table, so to speak, by i.i.d. samples for each $\Phi(\pi(s, a)) \in \{0, 10\}$. Unlike with two-player games, we need to also sample the actual state of the world, which we selected uniformly at random from the available states in the current decision problem. The expected fitness of phenotype t^i in decision problem D is then given by:

$$F_D(t^i) = \sum_{s \in S} \frac{\Phi(\pi(s, \hat{a}^i(u^i, \Gamma^i)))}{|S|},$$

with $\Gamma^i \subseteq \Delta(S)$. As phenotypes, we considered the original cast of four from Table 5.1, since altruistic and competitive types are meaningless in solitary decision situations. As before, the relevant fitness measure, defined in equation 5.3, was approximated by Monte Carlo simulations, the results of which are given in Table 5.5.

$$F(t^i) = \int F_D(t^i) dP_D(D) \tag{5.3}$$

Facts 5.1 and 5.2 still apply: (Reg, \tilde{P}) and (Mm, \tilde{P}) are behaviorally equivalent in general, and $(Reg, \tilde{\Gamma})$ is behaviorally equivalent to the former two in

¹¹Here we choose this formalization of decision problem to stay closer to the game-theoretic formulation of this and the previous chapter. A formalization closer to the chapter on rational choice would be to express a decision problem as a tuple $D = \langle S, X, A, \Phi \rangle$, where S, X and Φ are as before, and acts $a \in A$ are functions $a : S \rightarrow X$. It is however possible to translate one formulation into the other.

$(Reg, \tilde{\Gamma})$	$(Mm, \tilde{\Gamma})$	(Reg, \tilde{P})	(Mm, \tilde{P})
6.318	6.237	6.661	6.661

Table 5.5: Expected fitness approximated from 100.000 simulated solitary decision problems.

decision problems with two states and two acts. This shows in the results from table 5.5 in that the averages for (Reg, \tilde{P}) and (Mm, \tilde{P}) are identical. But since we included decision problems with more acts and more states as well, the average for regret minimizers $(Reg, \tilde{\Gamma})$ is not identical to the one of (Reg, \tilde{P}) and (Mm, \tilde{P}) . It is, in fact, lower, but again not as low as that of $(Mm, \tilde{\Gamma})$.

It follows that every relevant result we have seen about game situations is also borne out for solitary decisions. Evolutionary selection based on objective fitness will not select against regret-based choice principles, as these are indistinguishable from expected utility maximization when paired with precise beliefs. But when paired with imprecise beliefs, expected regret minimizers outperform expected utility maximizers. Consequently, if there is a chance, however small, that agents fall back on imprecise beliefs, evolution will actually positively select for regret-based decision criteria.

5.9 Conclusion

5.9.1 Subjective preferences or choice principles?

Before concluding, it is interesting to note that it may not always be straightforward to understand which individual components are the cause of different observed behaviors. Consider the example of the Hi-Lo fitness game for instance. An expected regret minimizer $(Reg, \Phi, \tilde{\Gamma})$ would choose action *II*. However, any expected utility maximizer $(Mm, \Phi, [\frac{1}{3} - \gamma, \frac{1}{3} + \delta])$, where $\frac{1}{3} - \gamma$ is the lower probability of the co-player playing *II* and $\frac{1}{3} + \delta$ is the upper probability, would also play action *II* for any $\gamma \in [0, \frac{1}{3}]$ and $\delta < 2\gamma$. In this case, observing action *II* can be explained in two different ways: the player might be either a maximally imprecise expected regret minimizer $(Reg, \Phi, \tilde{\Gamma})$, or a non-maximally imprecise expected utility maximizer $(Mm, \Phi, [\frac{1}{3} - \gamma, \frac{1}{3} + \delta])$, with $\gamma \in [0, \frac{1}{3}]$ and $\delta < 2\gamma$. Of course, if we knew the subjective belief of the player (because, for example, it comes from a learning process that we are fully informed about), then we could exclude one of the two cases and deduce the player's choice principle accordingly. We will follow this direction in the next chapter.

There is, however, an analogous but subtler issue that involves subjective preferences. Suppose that we define the following subjective utility:

- the *regret*: for all $G \in \mathcal{G}$,

$$\text{reg}(\pi(a, a')) := \Phi(\pi(a, a')) - \max_{a'' \in A} \Phi(\pi(a'', a')).$$

Then, the two types $(Reg, \Phi, \tilde{\Gamma})$ and $(Mm, \text{reg}, \tilde{\Gamma})$ would be behaviorally indistinguishable, for any $G \in \mathcal{G}$. In other words, simple regret minimizing behavior can be thought of as the result of two distinct processes: expected regret minimization acting on objective fitness (and maximally imprecise beliefs), or maxmin expected utility acting on regret utility (and maximally imprecise beliefs). In this chapter, opting for the first formalization was more a matter of intuition than a necessary choice, and there is no conclusive argument to prefer one to the other. Since the behavior of the two phenotypes is exactly the same, all the results presented above would equally hold if we substituted type $(Reg, \Phi, \tilde{\Gamma})$ with type $(Mm, \text{reg}, \tilde{\Gamma})$. At the same time, the conceptual interpretation of the results in Section 5.7 might change depending on the formalization we choose. In the first case, evolution favors one decision criterion over another; in the second, evolution selects a subjective preference over the other. While beliefs can be updated by learning and decision criteria can eventually be improved (perhaps) by studying some decision theory and game theory, subjective preferences (similarly to tastes) look like more rooted and innate characteristics of the individual. From a psychological point of view, acting at the level of beliefs and decision criteria seems very different than acting at the level of subjective preferences. We will expand on this issue in later chapters.

5.9.2 Precise and imprecise beliefs

So far, we have mainly focused on the evolutionary part of the story, and we have partially neglected the epistemic part, namely the ways of forming and updating beliefs about co-players' behavior or the actual state of the world. As mentioned earlier, there are motivations for this choice. The first is programmatic. Although objecting against keeping beliefs fixed is a reasonable point, it is in the spirit of evolutionary game theory (as we have seen in Section 3.2) to consider unsophisticated agents to begin with. The second reason is expositional. Evolutionary selection and belief updating are by and large distinct topics. For the completeness of our analysis we will have to bring them together, and allow the agents to learn about the population composition while playing the game of life. Nonetheless, to ease the exposition we decided to split the evolutionary part, that was the focus of this chapter, from the learning and belief updating, that will be the topic of the next chapter.

One last remark is in order here, before concluding. In line with the normative claims of Gilboa, Postlewaite and Schmeidler (see Chapter 1 and Section 2.2), we do not necessarily view the lack of a precise probabilistic belief as a cognitive limitation of the agents, but rather as the conscious awareness that in certain

situations a non-Bayesian belief is the best possible option given the amount and the quality of information available. Quoting [Gilboa and Marinacci, 2013] again, we agree that

Being able to admit ignorance is not a mistake. It is, we claim, more rational than to pretend that one knows what cannot be known.

In our population scenario, forcing the players to always have Bayesian beliefs about the co-player's actions, even when they do not know anything about the proportions of types in the population and they never met that co-player before, would be excessively demanding in our opinion. Similarly to Ellsberg's urn, it seems more reasonable to think that players (should) have imprecise beliefs about the behavior of a co-player randomly drawn from a population whose composition is unknown. The population context makes the emergence of imprecision in beliefs almost inevitable, especially when players from a big population have interacted only a few times, as in the quote from [Battigalli et al., 2015] above.

5.10 Appendix: Proofs

The proof of Proposition 5.4 relies on a partition of \mathcal{G} , and on some lemmas. For brevity, let us denote the regret minimizer $(Reg, \tilde{\Gamma})$ by R and the maximinimizer $(Reg, \bar{\Gamma})$ by M . Following equation (5.1), let $F_{\mathcal{G}}(X, Y)$ denote the expected fitness of type X against type Y on the possibly restricted class of fitness games \mathcal{G} .

Proof of Proposition 5.4. By definition of strict dominance, we have to show that in the class \mathcal{G} of symmetric 2×2 games with fitness numbers sampled from a set of i.i.d. values with at least 3 elements in the support, it holds that:

- (i) $F_{\mathcal{G}}(R, R) > F_{\mathcal{G}}(M, R)$;
- (ii) $F_{\mathcal{G}}(M, M) < F_{\mathcal{G}}(R, M)$.

A symmetric 2×2 game is fully specified by a quadruple of four numbers (a, b, c, d) .

	I	II
I	a	b
II	c	d

We partition the class \mathcal{G} as follows:

1. Coordination games \mathcal{C} : $\{(a, b, c, d) \in \mathcal{G} : a > c \text{ and } d > b\}$;
2. Anti-coordination games \mathcal{A} : $\{(a, b, c, d) \in \mathcal{G} : a < c \text{ and } d < b\}$;

3. Strong dominance games \mathcal{S} : $\{(a, b, c, d) \in \mathcal{G} : a > c \text{ and } b > d\} \cup \{(a, b, c, d) \in \mathcal{G} : a < c \text{ and } b < d\}$;
4. Weak dominance games \mathcal{W} : $\{(a, b, c, d) \in \mathcal{G} : a = c \text{ and } b \neq d\} \cup \{(a, b, c, d) \in \mathcal{G} : a \neq c \text{ and } b = d\}$;
5. Boring games \mathcal{B} : $\{(a, b, c, d) \in \mathcal{G} : a = c \text{ and } b = d\}$.

Before proving the lemmas, it is convenient to fix some notation. Let us call x, y, z the 3 elements in the support, and without loss of generality suppose that $x > y > z$. We denote by C a coordination game in \mathcal{C} with fitness a_C, b_C, c_C , and d_C ; similarly for games $A \in \mathcal{A}$, $S \in \mathcal{S}$, $W \in \mathcal{W}$, and $B \in \mathcal{B}$. Let us denote by I_{RC} the event that a R -player plays action I in the game C ; and similarly for action II , for a player of type M , and for games A, S, W and B . We first consider the case of i.i.d. sampling with discrete support.

Lemma 5.6. *R and M perform equally well in \mathcal{S} and in \mathcal{B} .*

Proof. By definition of regret minimization and maxmin it is easy to check that whenever in a game there is a strongly dominant action $a^{\$}$, then $a^{\$}$ is both the maxmin action and the regret minimizing action. Then, for all the games in \mathcal{S} , R chooses action a if and only if M chooses action a . Consequently, R and M always perform equally (well) in \mathcal{S} . In the case of \mathcal{B} it is trivial to see that all the players perform equally. \square

Lemma 5.7. *In \mathcal{W} , R strictly dominates M .*

Proof. Assume without loss of generality that $b = d$, and that $a > c$. There are two cases that we have to check: (i) $c < b = d$ and (ii) $c \geq b = d$. In the first case it is easy to see that R and M perform equally: act I is the choice of both R and M . In the case of (ii) instead we have that I is the regret minimizing action, whereas both actions have the same minimum and M plays $(\frac{1}{2}I; \frac{1}{2}II)$, since both I and II maximize the minimal fitness. Consider now a population of R and M playing games from the class \mathcal{W} . Whenever (i) is the case R and M perform equally well. But suppose $W \in \mathcal{W}$ and (ii) is the case. Then, $F_W(R, R) = a > \frac{1}{2}a + \frac{1}{2}c = F_W(M, R)$, whereas $F_W(M, M) = \frac{1}{4}a + \frac{1}{4}b + \frac{1}{4}c + \frac{1}{4}d < \frac{1}{2}a + \frac{1}{2}b = F_W(R, M)$. Hence, we have that in general $F_{\mathcal{W}}(R, R) > F_{\mathcal{W}}(M, R)$, and $F_{\mathcal{W}}(M, M) < F_{\mathcal{W}}(R, M)$. \square

Since it is not difficult to see that both (R, R) and (M, M) are strict Nash equilibria in \mathcal{C} , and that (R, R) and (M, M) are not Nash equilibria in \mathcal{A} , the main part of the proof will be to show that R strictly dominates C in the class $\mathcal{C} \cup \mathcal{A}$, that is:

- (i') $F_{\mathcal{C} \cup \mathcal{A}}(R, R) > F_{\mathcal{C} \cup \mathcal{A}}(M, R)$,
- (ii') $F_{\mathcal{C} \cup \mathcal{A}}(M, M) < F_{\mathcal{C} \cup \mathcal{A}}(R, M)$.

This part needs some more lemmas to be proven, but firstly we introduce the following bijective function ϕ between coordination and anti-coordination games.

Definition 5.8. ϕ . The permutation $\phi(a, b, c, d) = (c, d, a, b)$ defines a bijective function $\phi : \mathcal{C} \rightarrow \mathcal{A}$ that for each coordination game $C \in \mathcal{C}$ with fitness (a_C, b_C, c_C, d_C) gives the anti-coordination game $A \in \mathcal{A}$ with fitness $(a_A, b_A, c_A, d_A) = (c_C, d_C, a_C, b_C)$. Essentially, ϕ swaps rows in the fitness matrix.

Lemma 5.9. *Occurrence probability of C equals that of $\phi(C)$: $P(\phi(C)) = P(C)$.*

Proof. By definition, each game $C = (a_C, b_C, c_C, d_C)$ is such that $a_C > c_C$ and $d_C > b_C$, and each game $A = (a_A, b_A, c_A, d_A)$ is such that $a_A < c_A$ and $d_A < b_A$. Given that a, b, c, d are i.i.d. random variables and that a sequence of i.i.d. random variables is exchangeable, it is clear that the probability of (a_C, b_C, c_C, d_C) equals the probability of (c_C, d_C, a_C, b_C) . Hence, $P(\phi(C)) = P(C)$. \square

Lemma 5.10. *Let $P(E)$ be the probability of event E , e.g., $P(I_{RC})$ is the probability that a random R -player plays act I in coordination game C . It then holds that:*

- $P(I_{RC}) = P(II_{R\phi(C)})$, and $P(II_{RC}) = P(I_{R\phi(C)})$;
- $P(I_{MC}) = P(II_{M\phi(C)})$, and $P(II_{MC}) = P(I_{M\phi(C)})$.

Proof. It is easy to check that if $b_C - d_C > c_C - a_C$, a R -player plays action I in C ; that if $b_C - d_C < c_C - a_C$, R plays II ; and that if $b_C - d_C = c_C - a_C$, a R -player is indifferent between I and II in C , and so randomizes with $(\frac{1}{2}I; \frac{1}{2}II)$. Similarly, if $a_A - c_A > d_A - b_A$, a R -player plays action I in A ; if $a_A - c_A < d_A - b_A$, R plays II ; and if $a_A - c_A = d_A - b_A$, a R -player is indifferent between I and II in A , and randomizes with $(\frac{1}{2}I; \frac{1}{2}II)$. Consequently, if $b_C - d_C > c_C - a_C$, then $P(I_{RC}) = 1$, and by definition of ϕ we have $P(II_{R\phi(C)}) = 1$. Likewise, if $b_C - d_C < c_C - a_C$, then $P(II_{RC}) = 1 = P(I_{R\phi(C)})$; and if $b_C - d_C = c_C - a_C$, then $P(I_{RC}) = P(II_{RC}) = \frac{1}{2} = P(II_{R\phi(C)}) = P(I_{R\phi(C)})$.

In the same way, in coordination games we have that if $b_C > c_C$, a M -player plays I ; if $c_C > b_C$, a M -player plays II ; and if $b_C = c_C$, M is indifferent between I and II , and plays $(\frac{1}{2}I; \frac{1}{2}II)$. In anti-coordination games instead, if $a_A > d_A$, M plays I ; if $a_A < d_A$, M plays II ; if $a_A = d_A$, M plays $(\frac{1}{2}I; \frac{1}{2}II)$. By definition of ϕ : $P(I_{MC}) = 1 = P(II_{M\phi(C)})$ if $b_C > c_C$; $P(II_{MC}) = 1 = P(I_{M\phi(C)})$ if $c_C > b_C$; and $P(I_{MC}) = P(II_{MC}) = \frac{1}{2} = P(II_{M\phi(C)}) = P(I_{M\phi(C)})$ if $b_C = c_C$. \square

Lemma 5.11. *It holds that:*

- $a_C > d_C \rightarrow (I_{MC} \subseteq I_{RC})$;
- $a^C = d^C \rightarrow I_{MC} = I_{RC}$;

- $a_C < d_C \rightarrow (II_{MC} \subseteq II_{RC})$.

Proof. The event that R plays action I , I_{RC} , with positive probability is the event that $b_C - d_C \geq c_C - a_C$: if $b_C - d_C > c_C - a_C$, R plays I , and if $b_C - d_C = c_C - a_C$, R plays $(\frac{1}{2}I; \frac{1}{2}II)$. Similarly, the event that I_{MC} has positive occurrence is the event that $b_C \geq c_C$: if $b_C > c_C$, M plays I , and if $b_C = c_C$, M plays $(\frac{1}{2}I; \frac{1}{2}II)$. Then, I_{RC} implies that $b_C - d_C \geq c_C - a_C$, and I_{MC} implies that $b_C \geq c_C$. Moreover, on the assumption that $a_C > d_C$, it is easy to check that $b_C \geq c_C$ implies $b_C - d_C > c_C - a_C$. Hence, in any C with $a_C > d_C$ it holds that I_{MC} implies I_{RC} , i.e., $a_C > d_C \rightarrow (I_{MC} \subseteq I_{RC})$. Instead, it is possible that $a_C > d_C$, $b_C - d_C > c_C - a_C$ and $b_C < c_C$ hold simultaneously, so that $I_{MC} \not\subseteq I_{RC}$. By a symmetric argument it can be shown that $a_C < d_C \rightarrow (II_{MC} \subseteq II_{RC})$ too. Finally, when $a_C = d_C$ it holds that: $b_C - d_C > c_C - a_C$ iff $b_C > c_C$; $b_C - d_C < c_C - a_C$ iff $b_C < c_C$; and $b_C - d_C = c_C - a_C$ iff $b_C = c_C$. Hence, $a_C = d_C \rightarrow I_{MC} = I_{RC}$. \square

We are now ready to prove that $F_{C \cup A}(R, R) > F_{C \cup A}(M, R)$. With notation like $P(I_{RC} \cap I_{RC})$ denoting the probability that a random R -player plays I and another R -player plays I as well in game C , rewrite the inequality as:

$$\begin{aligned} & \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{RC}) \cdot a_C + P(II_{RC} \cap II_{RC}) \cdot d_C + P(I_{RC} \cap II_{RC}) \cdot b_C + P(II_{RC} \cap I_{RC}) \cdot c_C] \\ & + \sum_{A \in \mathcal{A}} P(A) [P(I_{RA} \cap I_{RA}) \cdot a_A + P(II_{RA} \cap II_{RA}) \cdot d_A + P(I_{RA} \cap II_{RA}) \cdot b_A + P(II_{RA} \cap I_{RA}) \cdot c_A] \\ & > \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{MC}) \cdot a_C + P(II_{RC} \cap II_{MC}) \cdot d_C + P(I_{RC} \cap II_{MC}) \cdot c_C + P(II_{RC} \cap I_{MC}) \cdot b_C] \\ & + \sum_{A \in \mathcal{A}} P(A) [P(I_{RA} \cap I_{MA}) \cdot a_A + P(II_{RA} \cap II_{MA}) \cdot d_A + P(I_{RA} \cap II_{MA}) \cdot c_A + P(II_{RA} \cap I_{MA}) \cdot b_A] \end{aligned}$$

By Lemma 5.9 and Lemma 5.10, we can express everything in terms of C only:

$$\begin{aligned} & \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{RC}) \cdot a_C + P(II_{RC} \cap II_{RC}) \cdot d_C + P(I_{RC} \cap II_{RC}) \cdot b_C + P(II_{RC} \cap I_{RC}) \cdot c_C \\ & + P(II_{RC} \cap II_{RC}) \cdot c_C + P(I_{RC} \cap I_{RC}) \cdot b_C + P(II_{RC} \cap I_{RC}) \cdot d_C + P(I_{RC} \cap II_{RC}) \cdot a_C] > \\ & \sum_{C \in \mathcal{C}} P(C) [P(I_{RC} \cap I_{MC}) \cdot a_C + P(II_{RC} \cap II_{MC}) \cdot d_C + P(I_{RC} \cap II_{MC}) \cdot c_C + P(II_{RC} \cap I_{MC}) \cdot b_C \\ & + P(II_{RC} \cap II_{MC}) \cdot c_C + P(I_{RC} \cap I_{MC}) \cdot b_C + P(II_{RC} \cap I_{MC}) \cdot a_C + P(I_{RC} \cap II_{MC}) \cdot d_C] \end{aligned}$$

This simplifies to:

$$\sum_C P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap II_{RC})) + b_C \cdot (P(I_{RC} \cap II_{RC}) + P(I_{RC} \cap I_{RC})) +$$

$$\begin{aligned}
& c_C \cdot (P(II_{RC} \cap I_{RC}) + P(II_{RC} \cap II_{RC})) + d_C \cdot (P(II_{RC} \cap II_{RC}) + P(II_{RC} \cap I_{RC})) \\
& > \sum_C P(C) [a_C \cdot (P(I_{RC} \cap I_{MC}) + P(II_{RC} \cap I_{MC})) + b_C \cdot (P(II_{RC} \cap I_{MC}) + P(I_{RC} \cap I_{MC})) \\
& + c_C \cdot (P(I_{RC} \cap II_{MC}) + P(II_{RC} \cap II_{MC})) + d_C \cdot (P(II_{RC} \cap II_{MC}) + P(I_{RC} \cap II_{MC}))]
\end{aligned}$$

Now let us split into $a > d$ and $a < d$, and consider $a > d$ first. Notice that, by Lemma 5.11, the case $a = d$ is irrelevant in order to discriminate between R and M . If $a > d$, by Lemma 5.11 we can eliminate the cases where R plays II and M plays I :

$$\begin{aligned}
& \sum_{C_{a>d}} P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap II_{RC})) + b_C \cdot (P(I_{RC} \cap II_{RC}) + P(I_{RC} \cap I_{RC})) + \\
& c_C \cdot (P(II_{RC} \cap I_{RC}) + P(II_{RC} \cap II_{RC})) + d_C \cdot (P(II_{RC} \cap II_{RC}) + P(II_{RC} \cap I_{RC}))] > \\
& \sum_{C_{a>d}} P(C) [a_C \cdot P(I_{RC} \cap I_{MC}) + b_C \cdot P(I_{RC} \cap I_{MC}) + c_C \cdot (P(I_{RC} \cap II_{MC}) + P(II_{RC} \cap II_{MC})) \\
& + d_C \cdot (P(II_{RC} \cap II_{MC}) + P(I_{RC} \cap II_{MC}))]
\end{aligned}$$

Rewrite:

$$\begin{aligned}
& \sum_{C_{a>d}} P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap II_{RC}) - P(I_{RC} \cap I_{MC})) \\
& + b_C \cdot (P(I_{RC} \cap II_{RC}) + P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC})) \\
& + c_C \cdot (P(II_{RC} \cap I_{RC}) + P(II_{RC} \cap II_{RC}) - P(I_{RC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) \\
& + d_C \cdot (P(II_{RC} \cap II_{RC}) + P(II_{RC} \cap I_{RC}) - P(II_{RC} \cap II_{MC}) - P(I_{RC} \cap II_{MC}))] \\
& > 0
\end{aligned}$$

We now distinguish between two cases: (1) $a - c = d - b$ and (2) $a - c \neq d - b$. Notice that $P(I_{RC} \cap II_{RC}) \neq 0$ if and only if case (1) obtains, and that $a > d$ and (1) imply II_{MC} . Then, from (1) we have:¹²

¹²Note that when we have only 3 elements in the support it is not guaranteed that case (1), together with $a > d$, may arise in a coordination game, whereas it is guaranteed that case (2), together with $a > d$, occurs with some positive probability. If we take for instance $x = 5, y = 2, z = 1$, then case (1) cannot obtain, whereas if we take $x = 3, y = 2, z = 1$, both (1) and (2) may obtain ($a = 3, b = 1, c = 2, d = 2$ for case (1), and $a = 3, b = 1, c = 2, d = 2$ for case (2)). Moreover, under the assumption that $a > d$, having 3 elements in the support is a necessary and sufficient condition for case (2) to have positive occurrence in a coordination game. As it will be clear in the following, a positive occurrence of case (2) only is enough for the theorem to hold.

$$\begin{aligned}
& \sum_{C_{a>d}} P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) + P(I_{RC} \cap II_{RC})) \\
& \quad + b_C \cdot (P(I_{RC} \cap II_{RC}) + P(I_{RC} \cap I_{RC})) \\
& \quad + c_C \cdot (P(II_{RC} \cap I_{RC}) + P(II_{RC} \cap II_{RC}) - P(I_{RC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) \\
& \quad + d_C \cdot (P(II_{RC} \cap II_{RC}) + P(II_{RC} \cap I_{RC}) - P(II_{RC} \cap II_{MC}) - P(I_{RC} \cap II_{MC}))] \\
& > 0
\end{aligned}$$

that is

$$\begin{aligned}
& \sum_{C_{a>d}} P(C) [a_C \cdot (\frac{1}{4} + \frac{1}{4}) + b_C \cdot (\frac{1}{4} + \frac{1}{4}) + c_C \cdot (\frac{1}{4} + \frac{1}{4} - \frac{1}{2} - \frac{1}{2}) + d_C \cdot (\frac{1}{4} + \frac{1}{4} - \frac{1}{2} - \frac{1}{2})] \\
& > 0
\end{aligned}$$

Since we have assumed $a - c = d - b$, the last inequality is not satisfied. We have instead:

$$\sum_{C_{a>d}} P(C) [\frac{1}{2}a_C + \frac{1}{2}b_C - \frac{1}{2}c_C - \frac{1}{2}d_C] = 0$$

This means that where $a_C > d_C$ and where (1) is the case, R and M are equally fit. This changes when we turn to (2). In that case, since $a_C > d_C \rightarrow (I_{MC} \subseteq I_{RC})$ by Lemma 5.11, we have that $P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC}) = P(I_{RC} \cap II_{MC})$. Moreover, if $a_C > d_C$, then $b_C \geq c_C$ implies $b_C - d_C > c_C - a_C$ (see Lemma 5.11). Consequently, when M plays either I or $(\frac{1}{2}I; \frac{1}{2}II)$, R always plays I . Hence, whenever $a_C > d_C$ and (2) obtain, it also holds that $P(II_{RC} \cap II_{MC}) = P(II_{RC} \cap II_{RC})$. In this case we can simplify:

$$\begin{aligned}
& \sum_{C_{a>d}} P(C) [a_C \cdot (P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC})) \\
& \quad + b_C \cdot (P(I_{RC} \cap I_{RC}) - P(I_{RC} \cap I_{MC})) \\
& \quad + c_C \cdot (P(II_{RC} \cap II_{RC}) - P(I_{RC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) \\
& \quad + d_C \cdot (P(II_{RC} \cap II_{RC}) - P(II_{RC} \cap II_{MC}) - P(I_{RC} \cap II_{MC}))] > 0
\end{aligned}$$

that is

$$\sum_{C_{a>d}} P(C) [P(I_{RC} \cap II_{MC}) \cdot (a_C + b_C - c_C - d_C)] > 0$$

We know that I_{RC} implies that $a_C - c_C \geq d_C - b_C$. Since we have assumed that $a_C - c_C \neq d_C - b_C$, we have that $a_C - c_C > d_C - b_C$. Hence, the inequality

$$\sum_{C_{a>d}} P(C) [P(I_{RC} \cap II_{MC}) \cdot (a_C + b_C - c_C - d_C)] > 0$$

is satisfied. So, when $a_C > d_C$, R strictly dominates M . Symmetrically, from $a < d$ and by distinguishing between the two cases (1) and (2) as before, in the end we get:

1. $\sum_{C_{a < d}} P(C) [-\frac{1}{2}a_C - \frac{1}{2}b_C + \frac{1}{2}c_C + \frac{1}{2}d_C] = 0;$
2. $\sum_{C_{a < d}} P(C) [P(II_{RC} \cap I_{MC}) \cdot (-a_C - b_C + c_C + d_C)] > 0.$

Hence, we can conclude that R strictly dominates M in the class $\mathcal{C} \cup \mathcal{A}$. Notice that in case of i.i.d. sampling with continuous support, games from \mathcal{B} and \mathcal{W} never arise, but the proof is the same for the remaining games in \mathcal{S} , \mathcal{C} and \mathcal{A} .

It remains to be shown that $F_{\mathcal{C} \cup \mathcal{A}}(M, M) < F_{\mathcal{C} \cup \mathcal{A}}(R, M)$. As before, spell this out as:

$$\begin{aligned} & \sum_C P(C) [P(I_{MC} \cap I_{MC}) \cdot a_C + P(II_{MC} \cap II_{MC}) \cdot d_C + P(I_{MC} \cap II_{MC}) \cdot b_C + P(II_{MC} \cap I_{MC}) \cdot c_C] \\ & + \sum_A P(A) [P(I_{MA} \cap I_{MA}) \cdot a_A + P(II_{MA} \cap II_{MA}) \cdot d_A + P(I_{MA} \cap II_{MA}) \cdot b_A + P(II_{MA} \cap I_{MA}) \cdot c_A] \\ & < \sum_C P(C) [P(I_{RC} \cap I_{MC}) \cdot a_C + P(II_{RC} \cap II_{MC}) \cdot d_C + P(I_{RC} \cap II_{MC}) \cdot b_C + P(II_{RC} \cap I_{MC}) \cdot c_C] \\ & + \sum_A P(A) [P(I_{RA} \cap I_{MA}) \cdot a_A + P(II_{RA} \cap II_{MA}) \cdot d_A + P(I_{RA} \cap II_{MA}) \cdot b_A + P(II_{RA} \cap I_{MA}) \cdot c_A] \end{aligned}$$

When $a > d$, similarly to the above derivation, we get:

$$\begin{aligned} & \sum_{C_{a > d}} P(C) [a_C \cdot (P(I_{MC} \cap I_{MC}) + P(I_{MC} \cap II_{MC}) - P(I_{RC} \cap I_{MC}) - P(I_{RC} \cap II_{MC})) \\ & + b_C \cdot (P(I_{MC} \cap I_{MC}) + P(I_{MC} \cap II_{MC}) - P(I_{RC} \cap I_{MC}) - P(I_{RC} \cap II_{MC})) \\ & + c_C \cdot (P(II_{MC} \cap I_{MC}) + P(II_{MC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) \\ & + d_C \cdot (P(II_{MC} \cap II_{MC}) + P(II_{MC} \cap I_{MC}) - P(II_{RC} \cap II_{MC}))] < 0 \end{aligned}$$

We now distinguish between (1) $b = c$, (2) $b > c$, and (3) $b < c$. Notice that $a > d$, combined with either (1) or (2), implies I_{RC} . Then we obtain:¹³

1. $\sum_{C_{a > d}} P(C) [-\frac{1}{2}a_C - \frac{1}{2}b_C + \frac{1}{2}c_C + \frac{1}{2}d_C] < 0;$
2. $\sum_{C_{a > d}} P(C) [a_C \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + b_C \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC}))] = 0;$
3. $\sum_{C_{a > d}} P(C) [a_C \cdot (-P(I_{RC} \cap II_{MC})) + b_C \cdot (-P(I_{RC} \cap II_{MC})) + c_C \cdot (P(II_{MC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) + d_C \cdot (P(II_{MC} \cap II_{MC}) - P(II_{RC} \cap II_{MC}))] \leq 0.$

When $a < d$, the derivation proceeds symmetrically and we get:

1. $\sum_{C_{a < d}} P(C) [\frac{1}{2}a_C + \frac{1}{2}b_C - \frac{1}{2}c_C - \frac{1}{2}d_C] < 0;$

¹³Note that here, when we only have 3 elements in the support, case (2) is impossible, but cases (1) and (3) occur with positive probability, and this is enough for our purpose.

2. $\sum_{C_{a < d}} P(C)[a_C \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + b_C \cdot (P(I_{MC} \cap I_{MC}) - P(I_{RC} \cap I_{MC})) + c_C \cdot (-P(II_{RC} \cap I_{MC})) + d_C \cdot (-P(II_{RC} \cap I_{MC}))] \leq 0;$
3. $\sum_{C_{a < d}} P(C)[c_C \cdot (P(II_{MC} \cap II_{MC}) - P(II_{RC} \cap II_{MC})) + d_C \cdot (P(II_{MC} \cap II_{MC}) - P(II_{RC} \cap II_{MC}))] = 0.$

Finally, we can conclude that $F_{C \cup A}(M, M) < F_{C \cup A}(R, M)$. As before, notice that, when we have i.i.d. sampling with continuous support, games from \mathcal{W} and \mathcal{B} never occur, but the proof is the same for all the other cases. Hence, both when the support of a, b, c, d is finite, and when the support is infinite, R strictly dominates M under the conditions assumed. \square

Chapter 6

Learning Players' Types

You can discover more about a person in an hour of play than in a year of conversation. (Plato)

In the preceding chapter we have studied the evolution of different subjective representations and decision criteria in the game of life, i.e., a multi-game framework where players are randomly matched with each other to play randomly selected two-player games. At the end of their life, players accumulate a certain amount of fitness (maybe money, or profit) that determines their offspring in the next generation.

However, along with the evolutionary dynamics, another fundamental dynamic may take place within the game of life if we endow the players with more sophisticated cognitive capabilities. Sophisticated players that recurrently interact with other members of the same population might also be able to learn statistical information about the population from past plays.

In classic evolutionary game theory, agents repeatedly playing the same game could eventually learn the proportions of different behaviors. Since in single-game models phenotypes usually coincide with expressed behavior, the only thing that players can learn is the frequency of different actions in their generation. The situation gets more complex once we look at the game of life introduced in the previous chapter, where a new fitness game is selected at each time. In the game of life, a phenotype does no longer coincide with a simple action, but represents instead a general behavior-generating mechanism that may select different actions in different games. The players' learning essentially aims at the proportions of these phenotypes in the population, and thus has to infer general behavior-generating mechanisms from observed behavior.

6.1 The Learning Model

This section explicitly specifies the dynamics according to which players learn in the game of life. We are interested in adding the learning part into the multi-game model of Chapter 5, and to study the evolutionary selection of choice principles also in the light of the learning dynamics.

As we have already seen, in standard evolutionary game theory, randomly selected agents from a given population recurrently play a unique, fixed fitness game,

$$G^0, G^0, G^0, G^0, \dots$$

The learning that we want to investigate instead takes place in the richer environment of the game of life, where the players play a sequence of randomly generated games,

$$G^1, G^2, G^3, G^4, \dots$$

In accordance with what we have presented in previous chapters, we postulate the following assumptions:

1. The games G^1, G^2, G^3, \dots are i.i.d. samples from some probability distribution P_G over the set of possible games \mathcal{G} .
2. The support of P_G is restricted to symmetric 2×2 games. These games can be parametrized in terms of four real numbers (a, b, c, d) , as illustrated in Table 6.1. Thus, P_G is effectively a probability distribution over \mathbb{R}^4 .
3. The four fitness values (a, b, c, d) are drawn i.i.d. from some probability measure P_V on the reals. The game distribution is thus the fourfold product of P_V , that is, $P_G = P_V^4$.
4. At each time $t \geq 1$, a symmetric 2×2 game G^t is sampled by drawing (a, b, c, d) , and two players are randomly selected from the population to play the game G^t .
5. At the end of the play, the two players split, another game G^{t+1} is sampled, other two players are randomly selected, and so on.
6. There is common knowledge among the players about the types in the population, but the proportions of types are unknown. For simplicity, we also assume that the outcome of each game is common knowledge within the population.
7. At the beginning of each new generation (i.e., at time $t = 0$), all players are born maximally uncertain about the proportions, and possibly reduce their uncertainty as they observe the sequence of plays in G^1, G^2, G^3, \dots in the subsequent time steps $t = 1, 2, 3, \dots$

	<i>I</i>	<i>II</i>
<i>I</i>	<i>a; a</i>	<i>b; c</i>
<i>II</i>	<i>c; b</i>	<i>d; d</i>

Table 6.1: General form table for symmetric 2×2 games

The present work has focused in particular on two alternatives for decision making: expected utility maximization and expected regret minimization. Here, we maintain this focus of interest by making a further assumption:

8. It is common knowledge that the types in the population are the expected maximinimizer (Mm, Φ, Γ) and the expected regret minimizer (Reg, Φ, Γ) defined in Chapter 5.

When we introduce the learning dynamics in the game of life, beliefs will no longer be a fixed, genetically predetermined trait of players, but will be instead subject to a process of repeated updating. Consequently, it becomes natural not to include any epistemic component Γ in the specification of a phenotype, and the beliefs needed by the decision criterion to produce a choice are endogenously determined through the learning process. Since, by point 8, the subjective utility u is fixed to the evolutionary fitness Φ for all players, a phenotype is now uniquely determined by the decision criterion \hat{a} . Consequently, the set of phenotypes in the population coincides with the two choice principles under consideration, that is, $\mathcal{T} = \{Mm, Reg\}$. From point 7, it thus follows that $\Gamma^0 = \Delta(\{Mm, Reg\})$: at time $t = 0$ players have maximal unmeasurable uncertainty about the proportions of types in the population. The first game G^1 is then selected, and, after observing the actions played in G^1 , players update their beliefs to $\Gamma^1 \subseteq \Delta(\{Mm, Reg\})$. Next, the second game G^2 is selected, and so on. For simplicity, we also assume that all players update their beliefs according to the same learning procedure. This assumption, together with points 6 and 7, implies that at any time t all the players in the population will hold the same belief Γ^t .

We remark again that, differently from Chapter 5, the beliefs of players are now expressed in terms of a (compact convex) set of probabilities over the possible *types* of the co-player, $\Gamma^t \subseteq \Delta(\{Mm, Reg\})$ for all $t \geq 0$, rather than in terms of a (compact convex) set of probabilities over the possible actions of the co-player. This reflects the target of the learning dynamics: players aim at learning the proportions of choice principles in the population, not only the frequency of single actions. We will later see that learning information on proportions of different types (of choice principles) translates into useful information on the co-player's future actions.

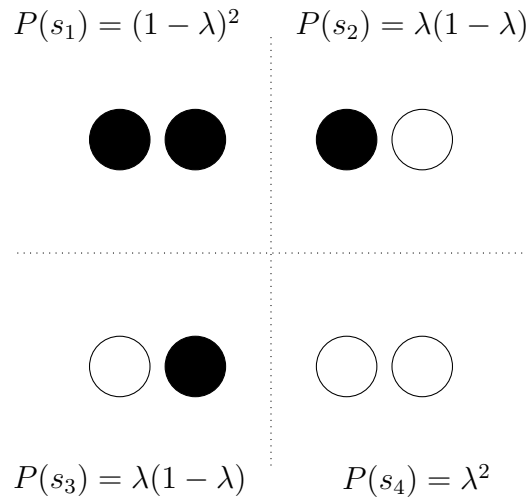


Figure 6.1: The outcome of two drawings under ambiguity.

6.2 Intermezzo: A Simple Extraction

To better understand the development of the chapter, the following example underlines the parallel with a decision-theoretic scenario by means of a simple extraction from an urn.

Consider an urn with 100 balls. Suppose that the only thing you know about the urn is that the balls are either white or black, but you don't have any further information about the relative proportions of the two colors. Suppose that your uncertainty is thus represented by the (compact and convex) set of probabilities $\tilde{\Gamma}$ that assigns lower probability of white $\underline{P}(W) = 0$ and upper probability of white $\overline{P}(W) = 1$, such that $\tilde{\Gamma} = \{P \in \Delta(\{W, B\}) : P(W) = \lambda, P(B) = 1 - P(W), \text{ for } \lambda \in [0, 1]\}$. As for the case of the random selection of players from a population with two types in Section 3.3, we can represent the possible distributions in the urn as an interval in $[0, 1]$, in particular $\tilde{\Gamma} = [0, 1]$, with the convention that each point $\lambda \in [0, 1]$ is a possible proportion of white balls. I will now draw a ball twice, with replacement, from the urn. Given the prior beliefs on white and black balls in the urn, your beliefs on the outcome of the two drawings must be as pictured by Figure 6.1, for $\lambda \in [0, 1]$.

The point of this example is that the obvious set of priors generated by the double extraction and depicted in Figure 6.1 is not convex, despite the fact that the belief $\tilde{\Gamma}$ over the white and black balls in the urn is convex. It is easy to see it, if we think of the set of priors over the possible outcome space of Figure 6.1 as a region in the three-dimensional simplex with vertices s_1, s_2, s_3 and s_4 (let us call this region $\tilde{\Gamma} \otimes \tilde{\Gamma}$). Since the probability $P(W) = 1$ is in $\tilde{\Gamma}$, then vertex s_4 of the three-dimensional simplex that gives probability 1 to a double white extraction is in $\tilde{\Gamma} \otimes \tilde{\Gamma}$ ($\lambda^2 = 1$ when $\lambda = 1$). Likewise, vertex s_1 giving probability 1 to two

black balls is also in $\tilde{\Gamma} \otimes \tilde{\Gamma}$, since $(1 - \lambda)^2 = 1$ when $\lambda = 0$. However, the convex combination $\frac{1}{2}s_1 + \frac{1}{2}s_4$ is not in $\tilde{\Gamma} \otimes \tilde{\Gamma}$ (since there is no possible $\lambda \in [0, 1]$ such that $\lambda^2 = (1 - \lambda)^2 = \frac{1}{2}$). Note, moreover, that if we condition on the event that the first ball is white W (or black B), the resulting set of probabilities is again convex, and it is precisely the original set $\tilde{\Gamma}$.

Suppose that you bet on a double white extraction from the urn, and that you will win \$10 if two white balls are drawn, so that $\pi(b_1, b_2) = \$10$ if $b_1 = b_2 = W$, and $\pi(b_1, b_2) = 0$ otherwise. Then, note that while, for each possible Bernoulli distribution P_λ with parameter λ , the expectation of a single white extraction

$$\mathbb{E}_{b_2 \sim P_\lambda} [\pi(W, b_2)]$$

is a linear function of the distribution P_λ ,

$$\mathbb{E}_{b_2 \sim P_\lambda} [\pi(W, b_2)] = \$10 \cdot \lambda + 0 \cdot (1 - \lambda),$$

the expectation on the double extraction, corresponding to a double integral, is not linear in the distribution P_λ :

$$\mathbb{E}_{b_1, b_2 \sim P_\lambda \times P_\lambda} [\pi(b_1, b_2)] = \$10 \cdot \lambda^2 + 0 \cdot (1 - \lambda)\lambda + 0 \cdot \lambda(1 - \lambda) + 0 \cdot (1 - \lambda)^2.$$

It is now the moment to reveal the mysterious exclamation marks of Table 3.3 in Section 3.3. The state space in Figure 3.2 is generated by a double extraction of players from the population, exactly as the outcome space of Figure 6.1 is the result of a double extraction of a ball from the urn. The set of probability distributions over the possible outcomes of the double extraction from the population is consequently not convex. [Walley, 1991] showed that given a compact and convex set of probabilities, minimal utilities are always attained at the extreme points of the set. Furthermore, the expectation for the case of a double extraction is not a linear function of the distribution P . In general, if $\pi(b_1, \dots, b_n)$ is a function of n Bernoulli variables, then the expectation is a polynomial of degree n in λ . The exclamation marks in Section 3.3 highlights the profiles where the minimal utility is not reached at the extreme points of the interval $[0.1, 0.9]$. This is possible because the expectations are no longer linear functions of the Bernoulli distribution P , and the set of product measures $\Gamma \otimes \Gamma$ is not convex.

6.3 The Game of Life with Incomplete Information and its Γ -Equilibria

Let us now get back to the game-theoretic context, and suppose that we are dealing with a given 2×2 symmetric game G^t and that two players have been randomly drawn from a population consisting of two possible types, $\mathcal{T} = \{t^1, t^2\}$. In the current framework, the two types are just two different decision criteria,

Mm and Reg , so that we can also call them *decision types*. Let us denote by $\lambda \in [0, 1]$ the relative proportion of type Reg in the population, so that $(1 - \lambda)$ is the proportion of type Mm . We argued in previous sections for taking into account non-probabilistic uncertainty about the proportions of types in such cases. In the model, players are born maximally uncertain at time $t = 0$, and at any time $t \geq 1$, they can narrow down the possible values of λ to some set $\Gamma^t = [\underline{\lambda}, \bar{\lambda}] \subseteq [0, 1]$. This gives rise to an incomplete information game under ambiguity. From Section 3.3 we know that in these cases the two solution concepts of ex ante and interim Γ -equilibrium do no longer coincide. It is thus natural to focus on the interim analysis of the game, and to consider the set of posterior probabilities Γ^t .¹

The set Γ^t leaves the players with ambiguity about the value of λ , and they will need to involve some decision criterion that can handle this non-probabilistic uncertainty. As we have seen in the previous chapter, there are different possible ways to deal with such situations, and to turn the overall uncertainty into an action choice. The way an agent resolves the uncertainty and gives rise to an order over her possible actions is encoded by her decision type.

Given a value $\lambda \in [0, 1]$, and a profile of policy functions $\vec{\sigma} = (\sigma_1, \sigma_2)$, the interim expected utility for type Mm of player i is

$$\begin{aligned} & \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\sigma_i, \sigma_{3-i}) | Mm] := \\ & (1 - \lambda)\Phi(\pi(\sigma_i(Mm), \sigma_{3-i}(Mm))) + \lambda\Phi(\pi(\sigma_i(Mm), \sigma_{3-i}(Reg))). \end{aligned}$$

Likewise, the interim expected utility for decision type Reg of player i is

$$\begin{aligned} & \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\sigma_i, \sigma_{3-i}) | Reg] := \\ & (1 - \lambda)\Phi(\pi(\sigma_i(Reg), \sigma_{3-i}(Mm))) + \lambda\Phi(\pi(\sigma_i(Reg), \sigma_{3-i}(Reg))) \\ & - \max_{a \in A} \{ (1 - \lambda)\Phi(\pi(a, \sigma_{3-i}(Mm))) + \lambda\Phi(\pi(a, \sigma_{3-i}(Reg))) \}. \end{aligned}$$

Hence, given a set $\Gamma^t \subseteq [0, 1]$ and a profile of policy functions $\vec{\sigma}^* = (\sigma_1^*, \sigma_2^*)$, the action $\sigma_i^*(Mm)$ is optimal for type Mm of player i if

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i: \mathcal{T} \rightarrow A} \min_{P_\lambda \times P_\lambda \in \Gamma \otimes \Gamma} \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\pi(\sigma_i, \sigma_{3-i}^*)) | Mm].$$

Note that a profile $\vec{\sigma}$, together with a belief Γ^t , determines a minimal and a maximal probability for the co-player playing II . Therefore, given a profile $\vec{\sigma}^*$,

¹There are at least two reasons why it is natural to decide for the interim perspective. The first is more technical: the sets Γ^t are convex and the interim expected utility is linear in the distributions P_λ , as observed in Section 6.2. As a consequence, the interim perspective relates more tightly to the representation results on ambiguity that we have seen in Chapter 2. The second reason is conceptual: in the population models of evolutionary game theory, players *are* types (or phenotypes) from the beginning to the end of their life, and there is no moment in time where they have to wait for a signal to be informed about the type they are. The players' perspective is always from the interim point of view.

we can rewrite the previous expression in the notation of Chapter 5 and say that action $\sigma_i^*(Mm)$ is optimal for type Mm if

$$\sigma_i^*(Mm) \in Mm(\Phi, \Gamma^t | \sigma_{3-i}^*),$$

where for $P_\lambda \in \Gamma^t$, $P_\lambda | \sigma_{3-i}^* \in \Delta(A)$ is defined by

$$(P_\lambda | \sigma_{3-i}^*)(a) = \sum_{t^i \in \mathcal{T}: \sigma_{3-i}^*(t^i) = a} P_\lambda(t^i),$$

and

$$\Gamma^t | \sigma_{3-i}^* := \{P_\lambda | \sigma_{3-i}^* \in \Delta(A) : P_\lambda \in \Gamma^t\}.$$

In a similar way, the action $\sigma_i^*(Reg)$ is optimal for type Reg of player i if

$$\sigma_i^*(Reg) \in Reg(\Phi, \Gamma^t | \sigma_{3-i}^*),$$

that is, if

$$\sigma_i^* \in \operatorname{argmax}_{\sigma: \mathcal{T} \rightarrow A} \min_{P_\lambda \times P_\lambda \in \Gamma \otimes \Gamma} \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\sigma, \sigma_{3-i}^*) | Reg].$$

The following definition then extends the notion of interim Γ -equilibrium to the case of a population where players differ in the decision criterion.

Definition 6.1. Interim Γ -equilibrium for decision types. In the incomplete information games under ambiguity that we are considering, a profile of policy functions $(\sigma_i^* : \{Mm, Reg\} \rightarrow A)_{i \in \{1,2\}}$ is an *interim Γ -equilibrium for decision types* if, for $i \in \{1, 2\}$,

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i: \mathcal{T} \rightarrow A} \min_{P_\lambda \times P_\lambda \in \Gamma \otimes \Gamma} \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\pi(\sigma_i, \sigma_{3-i}^*)) | Mm]$$

and

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i: \mathcal{T} \rightarrow A} \min_{P_\lambda \times P_\lambda \in \Gamma \otimes \Gamma} \mathbb{E}_{P_\lambda \times P_\lambda} [u_i(\sigma_i, \sigma_{3-i}^*) | Reg].$$

Whenever the set Γ^t is a singleton, $\Gamma^t = \{\lambda\}$, the previous definition reduces to the definition of interim equilibrium of the Bayesian game (see Section 3.1). But in this case, as we already know, the action that maximizes expected utility would be the same action that minimizes expected regret, so that the two decision criteria would always dictate the same action choice (Fact 5.1). The difference between the two types emerges only when players have *unmeasurable* uncertainty about the type of the co-player. In this sense, the game can only be either (the Bayesian elaboration² of) a game with complete information, or a game with incomplete information under ambiguity, without the “intermediate” stage of a standard Bayesian game with incomplete information.

²See [Battigalli, 2016].

Finally, Definition 6.1 can be generalized to situations where there are more decision criteria in the population, and also different subjective utilities u . In this case, the types in the population might not only differ in the decision criterion, but also with respect to the subjective utility that the criteria operate on. (A type would then be specified by a pair $t^j = (\hat{a}^j, u^j)$, as we have seen in the previous chapter.)

Definition 6.2. Interim Γ -equilibrium for the game of life. A profile of policy functions $\vec{\sigma}^* = (\sigma_1^*, \sigma_2^*)$ is defined to be an *interim Γ -equilibrium for the game of life* if it is the case for $i \in \{1, 2\}$ and all types $t^j = (\hat{a}^j, u^j) \in \mathcal{T}$ that

$$\sigma_i^*(t^j) \in \hat{a}^j(u^j, \Gamma^t | \sigma_{3-i}^*).$$

We say that a Γ -equilibrium $\vec{\sigma}^*$ is *symmetric* if all players use the same policy, $\sigma_1^* = \sigma_2^*$. Throughout most of this chapter we will only consider symmetric equilibria. Intuitively, this amounts to requiring that the action choices are stable with respect to types: all regret types *Reg* would choose the same action given the same game, and likewise for all maxmin types *Mm*. As we will see later, this assumption can be weakened, but it is a practical starting point for the analysis.

6.4 Symmetry and Learning

Consider the fitness game of Example 3.2 played in a population of *Mm* and *Reg* types where players have unmeasurable uncertainty about the relative shares of the two types, such that $\Gamma^t = [0.1, 0.9]$. Being an anti-coordination game, it is easy to check that neither of the two policies $\sigma^I = (I, I)$ and $\sigma^{II} = (II, II)$, where both types play the same action, can give rise to a symmetric Γ -equilibrium. Indeed, both maxmin and regret types have incentives not to coordinate in anti-coordination games. (A game that rewards coordination or anti-coordination does the same in terms of regret too.) We can consequently exclude the profiles $\vec{\sigma}^I$ and $\vec{\sigma}^{II}$ from the possible Γ -equilibria of the game. Similarly, it is straightforward to see that the profile $\vec{\sigma} = (\sigma^I, \sigma^{II}) = ((I, I), (II, II))$ where player 1 always plays action *I* and player 2 always plays *II*, independent of their types, is a Γ -equilibrium of the game (and the same would hold for the reversed equilibrium (σ^{II}, σ^I)). However, that equilibrium is not symmetric, and there are two reasons why it is not a good equilibrium for our analysis.

Firstly, the definition of the profile (σ^I, σ^{II}) entails the possibility of distinguishing the two players by role. Since we are dealing with symmetric fitness games and randomly selected members of a single population, this distinction is not easily accessible, especially from the players' perspective. In fact, the distinction between player 1 and player 2 is just a technical artifice to define the relevant equilibrium concepts: from an evolutionary point of view, the phenotypes are all that exists in the population. If we don't restrict our attention to symmetric

equilibria, the distinction between phenotypes would be overlooked in favor of the distinction between players, and it would become evolutionarily meaningless.

Secondly and more importantly, if we don't restrict the analysis to symmetric equilibria, the distinction between the two decision types would become meaningless also from a decision-theoretic point of view, and the possibility of learning would be prevented. Loosely speaking, in order to learn the proportions of decision types, players must be able to reason according to the logic: "if the player is of type *Mm*, she will choose action a' ; if the player is of type *Reg*, then she will choose action a ." If this reasoning cannot be carried out, then there is no possibility of associating players' types with observed behavior. Suppose for instance that, at step 4 of the procedure described in Section 6.1, four players, rather than two, are randomly selected and paired in two couples to play the game. For the sake of the example, suppose also that one couple consists of two types *Mm*, while the second couple of two types *Reg*. If we allowed the two couples to play asymmetric equilibria, for instance the equilibrium (σ^I, σ^{II}) , then the action choices would in no way reflect the difference in types: in both couples, we would have two players of the same type playing different actions. Refusing the restriction to symmetric Γ -equilibria would thus prevent the players from learning.

Note, en passant, that Chapter 5 already contains a possible analysis of such situations: if players are born with maximal ambiguity about the population and they can never learn, they will just keep the original set $\tilde{\Gamma}$ throughout their life, without ever reducing the uncertainty.

6.5 Strong Informativity

To ensure the possibility of learning, we are thus looking for games where the agents unambiguously reveal their type by playing. In that case, we would say that the game is *strongly informative*. But what are the strongly informative games then? Games with incomplete information under ambiguity might have more than one symmetric Γ -equilibrium. Consider for example the Hi-Lo fitness game of Table 3.2, and suppose that the players' uncertainty about the relative proportion λ of type *Reg* is represented by the set $\Gamma^t = [0.2, 0.6]$. We then have three symmetric interim Γ -equilibria: $((I, I), (I, I))$, $((II, II), (II, II))$, and $((I, II), (I, II))$. Indeed, when player 1 is of type *Mm*, the minimal expected utilities are

$$\begin{array}{ll} \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, I)) | Mm] = 1 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, I)) | Mm] = 0.1 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, I)) | Mm] = 0 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, I)) | Mm] = 0.8 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, II)) | Mm] = 0.4 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, II)) | Mm] = 0 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, II)) | Mm] = 0.2 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, II)) | Mm] = 2 \end{array}$$

where, to ease notation, we abbreviated the product measure $P_\lambda \times P_\lambda$ by the relevant parameter λ . When player 1 is of type *Reg* we have instead:

$$\begin{array}{ll} \min_\lambda \mathbb{E}_\lambda[u_1(I, (I, I)) | \text{Reg}] = 0 & \min_\lambda \mathbb{E}_\lambda[u_1(I, (II, I)) | \text{Reg}] = -1.7 \\ \min_\lambda \mathbb{E}_\lambda[u_1(II, (I, I)) | \text{Reg}] = -1 & \min_\lambda \mathbb{E}_\lambda[u_1(II, (II, I)) | \text{Reg}] = 0 \\ \min_\lambda \mathbb{E}_\lambda[u_1(I, (I, II)) | \text{Reg}] = -0.8 & \min_\lambda \mathbb{E}_\lambda[u_1(I, (II, II)) | \text{Reg}] = -2 \\ \min_\lambda \mathbb{E}_\lambda[u_1(II, (I, II)) | \text{Reg}] = -0.7 & \min_\lambda \mathbb{E}_\lambda[u_1(II, (II, II)) | \text{Reg}] = 0 \end{array}$$

Since the game and the players' uncertainty are symmetric, the minimal utilities are the same for player 2.

Hence, we are looking for games with a unique symmetric Γ -equilibrium where different types are associated with different actions. We thus call a Γ -equilibrium $\vec{\sigma}^*$ *revealing* if

1. $\vec{\sigma}^*$ is a strict symmetric Γ -equilibrium;
2. $\vec{\sigma}^*$ associates different types with different actions.

A game is said to be *strongly informative* when there is a revealing Γ -equilibrium, and that equilibrium is the unique symmetric Γ -equilibrium of the game.

For concreteness, consider again the fitness game of Example 3.2 played in a population of decision types *Mm* and *Reg*, where players have unmeasurable uncertainty such that $\underline{\lambda} = 0.4$ and $\bar{\lambda} = 0.8$. Being an anti-coordination game, the two symmetric profiles $\vec{\sigma}^I$ and $\vec{\sigma}^{II}$ cannot be Γ -equilibria, as we have noticed in the previous section. We are then left with two possible symmetric profiles, $((I, II), (I, II))$ and $((II, I), (II, I))$. It will be helpful to visualize them graphically.

Figure 6.2 pictures the situation for the case of $((I, II), (I, II))$. The graph on the left represents the expectations from the two actions for type *Mm*, while the one on the right corresponds to the expectations in terms of regret. In both graphs, the two vertical dashed lines delimit the player's expectation about the probability of action *II*, when $\underline{\lambda} = 0.4$ and $\bar{\lambda} = 0.8$ and the population is playing according to the policy function (I, II) , i.e., $\sigma_i(\text{Mm}) = I$ and $\sigma_i(\text{Reg}) = II$ for $i \in \{1, 2\}$. From the graph in Figure 6.2a, it is easy to see that action *II* has a lower minimal expectation within the dashed interval than action *I*, so that the best reply for a type *Mm* against a population playing according to the policy (I, II) would indeed be action *I*. The graph on the right instead shows that for type *Reg* the best reply to a population playing the policy (I, II) is action *II*. This implies that the profile $((I, II), (I, II))$ is a revealing Γ -equilibrium of the game. But is it the unique symmetric equilibrium? Neither the profile $\vec{\sigma}^I$ nor $\vec{\sigma}^{II}$ can be a Γ -equilibrium, since it is an anti-coordination game, but if the reverse revealing profile $((II, I), (II, I))$ is also a Γ -equilibrium, then the same type might play different actions, in accordance with different possible equilibria.

Figure 6.3 depicts the case of a population playing according to the policy (II, I) . Note that everything is as in Figure 6.2, apart from the position of the

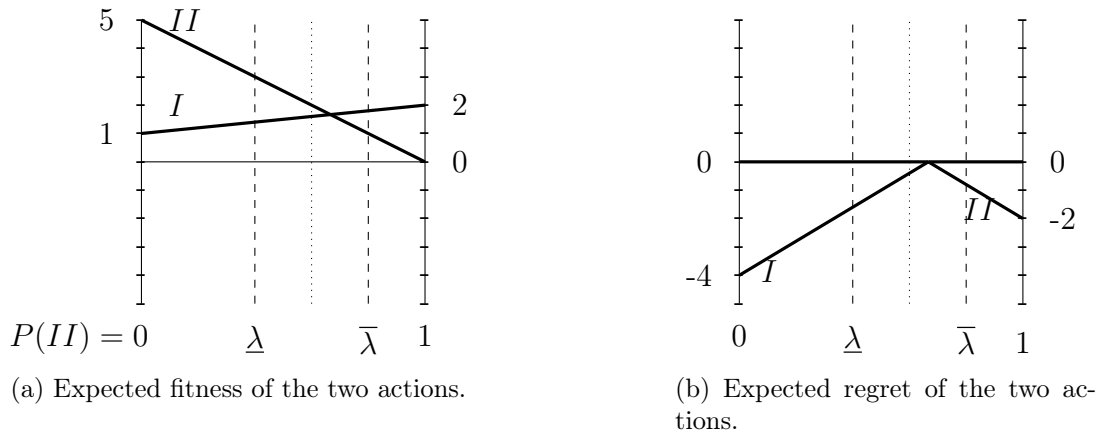


Figure 6.2: Maxmin and regret minimization for $\sigma_{3-i}(Mm) = I$ and $\sigma_{3-i}(Reg) = II$.




Figure 6.3: Maxmin and regret minimization for $\sigma_{3-i}(Mm) = II$ and $\sigma_{3-i}(Reg) = I$.

	$b < d$	$b = d$	$d < b$
$a < c$	II strongly dominant	II weakly dominant	Anti-coordination
$a = c$	II weakly dominant	Boring	I weakly dominant
$c < a$	Coordination	I weakly dominant	I strongly dominant

Table 6.2: A partition of the set of games according to how a compares to c , and how b compares to d .

two dashed lines: since we are now facing a population playing the policy (II, I) , we had to flip the two lines around to their mirror image. Looking at the minima of the two actions in the graph on the left, the minimal expected utility of action II is higher than that of action I , so that type Mm would choose to play II . In Figure 6.3b, the minimal utility of II is also higher than the minimal utility of I , and type Reg would also play II . Consequently, the profile $((II, I), (II, I))$ is not a Γ -equilibrium, and in this case the revealing Γ -equilibrium $((I, II), (I, II))$ is also the unique symmetric equilibrium of the game. Hence, the game is strongly informative. The numerical computation of minimal expected utilities confirms our analysis. When player 1 is of type Mm , the minimal utilities are

$$\begin{aligned}
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, I)) | Mm] &= 1 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, I)) | Mm] &= 1.2 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, I)) | Mm] &= 5 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, I)) | Mm] &= 2 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, II)) | Mm] &= 1.4 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, II)) | Mm] &= 2 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, II)) | Mm] &= 1 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, II)) | Mm] &= 0
\end{aligned}$$

When player 1 is of type Reg we have instead:

$$\begin{aligned}
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, I)) | Reg] &= -4 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, I)) | Reg] &= -2.8 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, I)) | Reg] &= 0 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, I)) | Reg] &= 0 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, II)) | Reg] &= -1.6 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, II)) | Reg] &= 0 \\
\min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, II)) | Reg] &= -0.8 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, II)) | Reg] &= -2
\end{aligned}$$

We thus found at least one strongly informative game for $\underline{\lambda} = 0.4$ and $\bar{\lambda} = 0.8$. However, our aim is to identify the conditions that can ensure the possibility of learning, independent of the uncertainty of the players and the composition of the population. To this end, the next lemmas generalize some ideas introduced in the previous examples.

Consider the partition of all possible symmetric 2×2 games (a, b, c, d) , according to how a compares to c , and how b compares to d , as shown in Table 6.2. We can then prove the following lemma.

Lemma 6.3. *Only anti-coordination games can be strongly informative.*

Proof. All proofs are in Section 6.10. □

As already noticed, none of the coordinating policies $\sigma^*(Mm) = \sigma^*(Reg) = a^*$ can be a Γ -equilibrium in an anti-coordination game, since both regret minimizing

and maximinimizing players will have an incentive to deviate from the coordinating policy functions in anti-coordination games. Hence, the strongly informative games are the anti-coordination games in which exactly one of the two anti-coordinating policies (i.e., such that either all types Mm play action I and all types Reg play action II , or vice versa) corresponds to the unique symmetric Γ -equilibrium of the game.

The next lemma expresses another property of strongly informative games, which narrows down this class even further.

Lemma 6.4. *An anti-coordination game can be strongly informative only if either $\underline{\lambda} < \frac{c-a}{c-a+b-d} < \bar{\lambda}$ or $1 - \bar{\lambda} < \frac{c-a}{c-a+b-d} < 1 - \underline{\lambda}$.*

In general, in any given population game with two types choosing from a menu of two actions, there are two symmetric and revealing policy profiles, since there are exactly two different ways for the two subpopulations to disagree. (If they had a menu of three actions, they could disagree in $3 \cdot 2$ different ways, with four, $4 \cdot 3$ different ways, and so on. For a menu of n actions they could disagree in $n^2 = n \cdot (n - 1)$ different ways.) Suppose, however, that both $\underline{\lambda} < \frac{c-a}{c-a+b-d} < \bar{\lambda}$ and $1 - \bar{\lambda} < \frac{c-a}{c-a+b-d} < 1 - \underline{\lambda}$ hold. Lemma 6.4 does not exclude that both revealing policies might correspond to a symmetric Γ -equilibrium. For strong informativity to obtain, we must have anti-coordination games such that only one between the policy $\sigma = (I, II)$ and the dual policy $\sigma = (II, I)$ defines a revealing Γ -equilibrium. The next lemma deals with this issue.

Lemma 6.5. *In any anti-coordination game, at most one between the two policies $\sigma = (I, II)$ and $\sigma = (II, I)$ can define a revealing Γ -equilibrium.*

6.6 The Set of Strongly Informative Games

It is not necessary, in general, that at least one of the revealing policy profiles corresponds to a Γ -equilibrium of the game. Whether this is the case depends on the fitness matrix of the game, and on the confidence interval $\Gamma^t = [\underline{\lambda}, \bar{\lambda}]$ that describes the players' current state of information. Since the fitness matrix of a particular game can be represented as a four-dimensional vector (a, b, c, d) , the set of strongly informative games corresponds, for a fixed $\Gamma = [\underline{\lambda}, \bar{\lambda}]$, to a certain region in the four-dimensional space. In this section, we will give a characterization of this region and provide conditions on the distribution over fitness matrices under which games fall inside of this strongly informative region.

Lemma 6.3 tells us that only anti-coordination games can be strongly informative, and Lemma 6.4 puts more restrictive necessary conditions for strong informativity to obtain. As a consequence, not every anti-coordination game is strongly informative. For example, it is possible that, for an interval $[\underline{\lambda}, \bar{\lambda}]$, an anti-coordination game (a, b, c, d) is such that neither $\underline{\lambda} < \frac{c-a}{c-a+b-d} < \bar{\lambda}$ nor

$1 - \bar{\lambda} < \frac{c-a}{c-a+b-d} < 1 - \underline{\lambda}$ holds. Moreover, it is also possible that one or both of $\underline{\lambda} < \frac{c-a}{c-a+b-d} < \bar{\lambda}$ and $1 - \bar{\lambda} < \frac{c-a}{c-a+b-d} < 1 - \underline{\lambda}$ hold, but still the same action is the best reply for both types.

Suppose therefore that an anti-coordination game (a, b, c, d) is given, and suppose that the population is split into a subpopulation of types Mm playing action I , and a subpopulation of types Reg playing action II . Conditional on player i being of type Mm , action I is strictly optimal, namely a Γ -best reply, if

$$\min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((I, II), (I, II))) | Mm] > \min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((II, II), (I, II))) | Mm].$$

Given the actions that are being played in the two subpopulations, this amounts to saying that I is strictly optimal if

$$\min_{\lambda} \{(1 - \lambda)a + \lambda b\} > \min_{\lambda} \{(1 - \lambda)c + \lambda d\}.$$

Since the minimum of a linear function over a closed interval is attained at one of the endpoints, this is equivalent to

$$\min \{(1 - \underline{\lambda})a + \underline{\lambda}b, (1 - \bar{\lambda})a + \bar{\lambda}b\} > \min \{(1 - \underline{\lambda})c + \underline{\lambda}d, (1 - \bar{\lambda})c + \bar{\lambda}d\}.$$

This can be written a bit more simply in terms of the following quantities (see Figure 6.4):

$$\begin{aligned} a' &:= (1 - \underline{\lambda})a + \underline{\lambda}b = a + \underline{\lambda}(b - a) \\ b' &:= (1 - \bar{\lambda})a + \bar{\lambda}b = a + \bar{\lambda}(b - a) \\ c' &:= (1 - \underline{\lambda})c + \underline{\lambda}d = c + \underline{\lambda}(d - c) \\ d' &:= (1 - \bar{\lambda})c + \bar{\lambda}d = c + \bar{\lambda}(d - c) \end{aligned} \tag{6.1}$$

Using this notation, the condition under which it is strictly optimal for types Mm to play action I against regret minimizers playing action II is that

$$\min \{a', b'\} > \min \{c', d'\}.$$

By definition, (a, b, c, d) is an anti-coordination game if and only if

$$c > a \quad \text{and} \quad b > d.$$

For $\underline{\lambda} < \frac{c-a}{c-a+d-b} < \bar{\lambda}$ (that is, Γ is nonempty and the necessary condition of Lemma 6.4 is satisfied), this entails

$$c' > a' \quad \text{and} \quad b' > d'.$$

Under these assumptions, we can make the following observations:

- If $a' < d'$, then a' is smaller than both c' and d' ; it follows that

$$\min \{a', b'\} < \min \{c', d'\},$$

and types Mm are not using a best reply.

- If $a' > d'$, then d' is smaller than both a' and b' ; it follows that

$$\min \{a', b'\} > \min \{c', d'\},$$

and types Mm are using a best reply.

- If $a' = d'$, then

$$\min \{a', b'\} = \min \{c', d'\} = a' = d',$$

and types Mm would be indifferent between the two actions.

When the population is split according to policy (I, II) , action I is thus strictly optimal for types Mm if and only if $a' > d'$. Expanding the definition of a' and d' , this can be reformulated as

$$a + \underline{\lambda}(b - a) > c + \bar{\lambda}(d - c),$$

which is equivalent to

$$(1 - \bar{\lambda})c + \bar{\lambda}d < (1 - \underline{\lambda})a + \underline{\lambda}b. \quad (6.2)$$

When we state the condition in terms of the last inequality, it is evident that it expresses a relation between the height of the line corresponding to action I at $\lambda = \underline{\lambda}$ (given by the convex combination $(1 - \underline{\lambda})a + \underline{\lambda}b$) and the height of the line corresponding to action II at $\lambda = \bar{\lambda}$ (given by the convex combination $(1 - \bar{\lambda})c + \bar{\lambda}d$). This provides the necessary and sufficient condition under which action I is the unique best reply for a player of type Mm when the population is playing according to the policy (I, II) .


Consider now the subpopulation of regret minimizers. In an anti-coordination game, these players are responding optimally to the population playing the policy (I, II) by choosing action II if and only if

$$c' - a' > b' - d'.$$

Expanding the definitions and rearranging the terms, this can be reformulated as

$$\frac{c - a}{c - a + b - d} > \frac{\underline{\lambda} + \bar{\lambda}}{2}. \quad (6.3)$$

Hence, for a given anti-coordination game, the policy function (I, II) , in which maxmin types play action I and regret minimizers play action II , defines a symmetric Γ -equilibrium when the conditions 6.2 and 6.3 are simultaneously satisfied. Therefore, when the conditions 6.2 and 6.3 hold simultaneously, the game is strongly informative.



(a) A strongly informative game where the population plays according to the policy (II, I) .

(b) A strongly informative game where the population plays according to the policy (I, II) .

Figure 6.4: Examples of strong informativity in anti-coordination games.

If we consider the case in which the subpopulation of types Mm plays II and the subpopulation of types Reg plays I , then by the dual argument the quantities a', b', c', d' are defined as

$$\begin{aligned}
 a' &:= b + \bar{\lambda}(a - b) \\
 b' &:= b + \underline{\lambda}(a - b) \\
 c' &:= d + \bar{\lambda}(c - d) \\
 d' &:= d + \underline{\lambda}(c - d)
 \end{aligned} \tag{6.4}$$

Accordingly, the necessary and sufficient condition for action II to be the unique best reply of Mm is given by

$$\underline{\lambda}c + (1 - \underline{\lambda})d > \bar{\lambda}a + (1 - \bar{\lambda})b, \tag{6.5}$$

whereas the necessary and sufficient condition for types Reg to play I is

$$\frac{c - a}{c - a + b - d} < \frac{(1 - \bar{\lambda}) + (1 - \underline{\lambda})}{2}. \tag{6.6}$$

Figure 6.4 depicts two games that exemplify these two possibilities for strong informativity in anti-coordination games.

In order to summarize these results in a more convenient way and to prove the next proposition, let us abbreviate the denominator by

$$Z := c - a + b - d$$

and notice that in Figures 6.4a and 6.4b, the intersection of the lines corresponding to action I and action II lies precisely at $\frac{c-a}{Z}$. From these premises it is possible to show that the conditions ensuring that, for any interval $[\underline{\lambda}, \bar{\lambda}]$ with $\underline{\lambda} \neq \bar{\lambda}$, the set of strongly informative games is nonempty (and, hence, that players can learn) are given in the following proposition.

Proposition 6.6. *The distribution λ^* of regret types in the population is learnable if and only if, for any $\underline{\lambda} \in [0, 1)$ and $\bar{\lambda} \in (\underline{\lambda}, 1]$, there is positive probability of having anti-coordination games (a, b, c, d) such that at least one of the following two conditions is satisfied:*

1. $\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c-a}{Z} < \frac{\bar{\lambda}(c-d) + \underline{\lambda}(b-a)}{Z}$;
2. $\frac{(1-\underline{\lambda})(c-d) + (1-\bar{\lambda})(b-a)}{Z} < \frac{c-a}{Z} < \frac{(1-\bar{\lambda}) + (1-\underline{\lambda})}{2}$.

Proof. In appendix 6.10. □

Although Proposition 6.6 may seem rather cryptic and technical, we will see that it has a clean geometrical representation, and that its conditions are satisfied in many simple scenarios.

Notice that, for the first condition to hold, it is necessary that

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{\bar{\lambda}(c-d) + \underline{\lambda}(b-a)}{Z}.$$

It turns out that this equation reduces to

$$b - a < c - d.$$

Symmetrically, the second condition can obtain only if

$$\frac{(1-\underline{\lambda})(c-d) + (1-\bar{\lambda})(b-a)}{Z} < \frac{(1-\bar{\lambda}) + (1-\underline{\lambda})}{2},$$

which is equivalent to

$$b - a > c - d.$$

Consequently, an anti-coordination game (a, b, c, d) such that $\frac{b-a}{Z} = \frac{1}{2}$ is never strongly informative. Figure 6.5b represents an example of this case. In general, notice that the quantity $\frac{b-a}{Z}$ distinguishes between the two possible informativities of the game: $\frac{b-a}{Z} < \frac{1}{2}$ is equivalent to $b - a < c - d$, and $\frac{b-a}{Z} > \frac{1}{2}$ is equivalent to $b - a > c - d$. Figure 6.4a is an example of $1 < \frac{b-a}{Z}$, while Figure 6.4b is an example of $\frac{b-a}{Z} < 0$. We can see from the figures that in both cases the two actions are represented by lines which are both increasing (Figure 6.4a) or both decreasing (Figure 6.4b). When $0 < \frac{b-a}{Z} < 1$ instead, the slopes of the two lines have different signs (one positive and one negative), as in Figure 6.5.

It is then possible to visualize the space of strongly informative games, for a given interval $[\underline{\lambda}, \bar{\lambda}]$, in two-dimensional space (Figure 6.6). The horizontal dotted lines represent the intervals $[\underline{\lambda}, \bar{\lambda}]$ and $[1 - \bar{\lambda}, 1 - \underline{\lambda}]$. We can see that the space of informative games is bounded both from above and from below by these lines, such that only anti-coordination games (a, b, c, d) with $\underline{\lambda} < \frac{c-a}{Z} < \bar{\lambda}$ or $1 - \bar{\lambda} < \frac{c-a}{Z} < 1 - \underline{\lambda}$ can be strongly informative. This corresponds to the necessary




Figure 6.5: Examples of strictly informative and uninformative games.




Figure 6.6: A representation of informativity in anti-coordination games.

condition of Lemma 6.4. The white area corresponds to informative games where *Mm* plays *I* and *Reg* plays *II*, while the black area corresponds to informative games where the population plays according to the dual policy (*II*, *I*). There is no region of space where the two colors overlap. As proven in Lemma 6.5, there is no anti-coordination game that can be informative in two different ways, i.e. such that the two policies (*I*, *II*) and (*II*, *I*) define two different revealing Γ -equilibria. It is evident from the picture that all informative games with policy (*I*, *II*) are on the left of $\frac{b-a}{Z} = \frac{1}{2}$, and all informative games with the dual policy (*II*, *I*) lie on the right of $\frac{b-a}{Z} = \frac{1}{2}$. Finally, notice that the line passing through the points $(0, \bar{\lambda})$ and $(1, \underline{\lambda})$ has equation

$$y = \bar{\lambda} + (\underline{\lambda} - \bar{\lambda})x,$$

that is,

$$\frac{c-a}{Z} = \bar{\lambda} + (\underline{\lambda} - \bar{\lambda})\frac{b-a}{Z},$$

which in turn reduces to

$$\frac{c-a}{Z} = \frac{\bar{\lambda}(c-d) + \underline{\lambda}(b-a)}{Z},$$

that is exactly one of the relevant edges of the first condition in Proposition 6.6. By symmetry, the same holds for the line passing through the points $(0, 1 - \underline{\lambda})$ and $(1, 1 - \bar{\lambda})$ and the second condition of Proposition 6.6.

6.7 Asymmetric Equilibria and Behavioral Assumptions

So far, we took into account anti-coordination games only. The reason was that, if we restrict our attention to symmetric equilibria, (some of) those games had the advantage of admitting a unique revealing Γ -equilibrium. As argued before, allowing for asymmetric equilibria may look like a departure from the decision-theoretic and evolutionary intuitions about types, but, after all, from a purely game-theoretic point of view one might object that asymmetric equilibria should be considered as accessible as symmetric equilibria, and that there is no cogent game-theoretic reason to exclude asymmetric equilibria from the analysis. This section is meant to offer an alternative reply to this possible objection.

It is evident that if we just included asymmetric equilibria in the analysis, we would be left with no strongly informative games at all. Indeed, their strong informativity stemmed from the exclusion of asymmetric equilibria: in certain anti-coordination games, the revealing equilibrium was the unique equilibrium available precisely because we prevented players from playing asymmetric equilibria. According to Lemma 6.5, however, any anti-coordination game admits at

most one revealing Γ -equilibrium, plus the two uninformative asymmetric equilibria $((I, I), (II, II))$ and $((II, II), (I, I))$. A first important step for the development of this section is that the equivalent of Lemma 6.5 also holds for coordination games.

Lemma 6.7. *In any coordination game, at most one between the two policies $\sigma = (I, II)$ and $\sigma = (II, I)$ can define a revealing Γ -equilibrium.*

Proof. Analogous to the proof of Lemma 6.5. □

If we want to drop the restriction to symmetric equilibria, it is helpful to start off by thinking what actions the two types would choose when faced with a coordination game.

Coordination games were immediately discarded from the analysis, inasmuch as the two uninformative profiles, $((I, I), (I, I))$ and $((II, II), (II, II))$, are always symmetric Γ -equilibria for all coordination games. On the other hand, the fitness achieved in coordination games is crucial for the evolutionary results of Chapter 5. In order to extend the arguments of Chapter 5, the behavior in coordination games must also be taken into consideration (as well as the behavior in uninformative anti-coordination games).

To fix ideas, it may be useful to start with a simple example. Suppose that a new generation was just born (so that players are still maximally uncertain about the distribution of types, $\Gamma^0 = [0, 1]$), and the first fitness game that is played is the same that we have seen at the end of Section 5.6:

	I	II
I	2; 2	1; 0
II	0; 1	5; 5

When we presented simple fitness maximization $Mm(\Phi, \tilde{\Gamma})$ and simple regret minimization $Reg(\Phi, \tilde{\Gamma})$ there, we saw that simple maxmin $Mm(\Phi, \tilde{\Gamma})$ would dictate to choose action I , while simple regret minimization $Reg(\Phi, \tilde{\Gamma})$ would suggest action II . According to the analysis of Chapter 6, where the primary source of uncertainty was on the co-player's type rather than directly on the co-player's actions, there are three interim symmetric Γ -equilibria for $\Gamma = [0, 1]$: $((I, I), (I, I))$, $((II, II), (II, II))$, and $((I, II), (I, II))$. Indeed, when player 1 is of type Mm , the minimal expected utilities are:

$$\begin{array}{ll}
 \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, I)) | Mm] = 2 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, I)) | Mm] = 1 \\
 \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, I)) | Mm] = 0 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, I)) | Mm] = 0 \\
 \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, II)) | Mm] = 1 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, II)) | Mm] = 1 \\
 \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, II)) | Mm] = 0 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, II)) | Mm] = 5
 \end{array}$$

When player 1 is of type *Reg* we have instead:

$$\begin{aligned} \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, I)) | \text{Reg}] &= 0 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, I)) | \text{Reg}] &= -4 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, I)) | \text{Reg}] &= -2 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, I)) | \text{Reg}] &= -2 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (I, II)) | \text{Reg}] &= -4 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(I, (II, II)) | \text{Reg}] &= -4 \\ \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (I, II)) | \text{Reg}] &= -2 & \min_{\lambda} \mathbb{E}_{\lambda}[u_1(II, (II, II)) | \text{Reg}] &= 0 \end{aligned}$$

Since the game is symmetric, the utilities of Player 2 are identical to those of Player 1.

When the players' uncertainty is lifted from the direct behavior of the co-player to the type of the co-player, such situations arise, and the action choices are no longer determined in a unique and obvious way, given the multiplicity of Γ -equilibria. This problem may noticeably echo the old issue of equilibrium selection (see [Harsanyi and Selten, 1988]). There seems to be a tension here between the analysis of Chapter 5 and the equilibrium analysis of games with incomplete information under ambiguity of this chapter. To assign behavior to types in each game that is played and to compute the accumulated fitness at the end of the generation, we have to resolve the tension between the two analyses and to decide how different types will behave when facing situations like that of the last example.

In our opinion, in front of multiple (symmetric or asymmetric) equilibria, a simple equilibrium analysis assuming that players always end up in one of the equilibria would be unsatisfactory. Backing the equilibrium play by standard evolutionary arguments based on single-game models would also not be maintainable given the current dynamics: since the games change continuously, there is no time for evolution to select an equilibrium among the many in each single game. How could players always end up in an equilibrium, even in situations where equilibria are multiple and the game is played only once and never repeated again? When looking for reasonable behavioral assumptions, we thus prefer to focus on the notion of *rationalizable* play, rather than on equilibrium play. The behavioral assumptions that we postulate are then the following.

Assumption 6.8. (Common belief in) Rationality. *In every game, each type plays a rationalizable action.*

Assumption 6.9. Type's preference. *In games with a revealing Γ -equilibrium, each type prefers the action specified by the revealing equilibrium. In games with no revealing Γ -equilibrium, each type plays one of the rationalizable actions at random.*

Assumption 6.8 just puts rationality constraints on possible action choices: roughly speaking, if an action cannot be justified as rational by any reasonable conjecture about the co-player, then that action will not be chosen. The second

assumption requires something more than (common belief in) rationality, and it produces a ranking over rationalizable actions.

To better understand the rationale behind Assumption 6.9, it may be useful to consider again the example at the beginning of this section. In that game, both types (and both players) have incentives to coordinate with the action of the co-player, but there is no obvious way to jointly decide for one action over the other. Ultimately, however, players have to play the game, and therefore they have to choose an action. Notice then that in the previous example action I is optimal for types Mm independent of the choice of types Reg : if types Reg play action I , types Mm should choose action I in order to best reply; if types Reg play action II , both actions are optimal for types Mm . Conversely, the same is true for action II and types Reg . This is the intuition behind the preference of a type for one rationalizable action over the other in the presence of a revealing Γ -equilibrium. In turn, the presence and the uniqueness of revealing Γ -equilibria (Lemma 6.5 and Lemma 6.7) comes from and reflects the decision-theoretic foundations of different types: in the previous example the reverse revealing behavior (II, I) would be at odds with the interpretation of the two types as different decision criteria. The fact that $((II, I), (II, I))$ is never a Γ -equilibrium of the game mirrors the difference in choice principles between the two types.

The introduction of the two behavioral assumptions allows to drop the restriction to symmetric equilibria (since revealing behavior is now derived from rationalizability and types' preferences), and to include also coordination games in the class of informative games.

Specifically, by the same argument as for the case of anti-coordination games, we obtain that the necessary and sufficient condition for I to be a strict best reply for Mm in a coordination game (a, b, c, d) against a population playing according to policy (I, II) is

$$(1 - \bar{\lambda})a + \bar{\lambda}b > (1 - \underline{\lambda})c + \underline{\lambda}d. \quad (6.7)$$

Likewise, a type Reg is responding optimally to the population playing the policy (I, II) by choosing action II if and only if

$$\frac{c - a}{c - a + b - d} < \frac{\lambda + \bar{\lambda}}{2}. \quad (6.8)$$

The space of informative games is therefore extended to all coordination games satisfying equations 6.7 and 6.8. So, for a given interval $\Gamma = [\underline{\lambda}, \bar{\lambda}]$, the space of informative coordination games is shown in Figure 6.7. Proposition 6.6 is thus extended to include the case of coordination games as well.

Proposition 6.10. *Given Assumption 6.8 and Assumption 6.9, the distribution λ^* of regret types in the population is learnable if and only if Proposition 6.6 holds, or, for any $\underline{\lambda} \in [0, 1)$ and $\bar{\lambda} \in (\underline{\lambda}, 1]$, there is positive probability of having*

The following proposition shows that learnability can be guaranteed under rather general circumstances. In particular, if fitness values are given by i.i.d drawings of a continuous random variable, the conditions of Proposition 6.6 and Proposition 6.10 follow and learning is ensured.

Corollary 6.11. *When games come from i.i.d. drawings of a continuous random variable, then the true distribution λ^* is learnable.*

6.9 Conclusion

Players in the game of life are now allowed the possibility of learning. We provided the formal conditions for the distribution of the decision types in the population to be learnable, and we showed how it is possible to infer behavior-generating mechanisms from observed behavior. To this end, we had to define suitable equilibrium concepts for incomplete information games under ambiguity, taking into account the variability in the players' decision criteria too. Overall, the results of this chapter might look mostly technical, but they guarantee the possibility of learning for rather intuitive scenarios, e.g., when fitness values are given by i.i.d. drawings of a continuous random variable, and exclude perfect learnability in other simple circumstances, for instance whenever fitness values are discrete.

Through the introduction of learning, modeled as process that reduces the ambiguity in front of incoming evidence, we also want to offer a formal theory of belief formation in the framework of the game of life. In line with the position of Gilboa, Postlewaite and Schmeidler presented in Chapter 1, we aim in this way at backing the project of having beliefs justified by evidence, as a basis for “a systematic way of predicting which beliefs agents might hold in various environments.”³

6.10 Appendix: Proofs

The following is the proof of Lemma 6.3.

Proof. A symmetric 2×2 game can be represented as a payoff matrix with entries (a, b, c, d) , so the set of possible games can be identified with (a subset of) \mathbb{R}^4 . We divide this space up into nine regions according to how a compares to c , and how b compares to d . As shown in Table 6.2, each of these possibilities corresponds to a certain game-theoretic property. Note that all of these properties are preserved under the regret transformation (e.g., coordination games are coordination games for types Mm as well as for types Reg).

By definition, a game is not strongly informative if there is a policy function $\sigma^*(Mm) = \sigma^*(Reg) = a^*$ such that the profile $\vec{\sigma}^*$ is a Γ -equilibrium. However,

³[Gilboa et al., 2012]

whenever some action a^* is weakly or strongly dominant, the policy function $\sigma^*(Mm) = \sigma^*(Reg) = a^*$ defines a symmetric Γ -equilibrium. We can thus immediately discard seven of these nine cases as uninformative.

This leaves us with coordination and anti-coordination games. However, in a coordination game, the coordinating policy function $\sigma^*(Mm) = \sigma^*(Reg) = a^*$ defines a symmetric Γ -equilibrium for $a^* \in \{I, II\}$. Hence, coordination games are never strongly informative. It follows that if there exist any strongly informative games at all, they must be anti-coordination games. \square

The following is the proof of Lemma 6.4.

Proof. Suppose that $\underline{\lambda} < \frac{c-a}{c-a+b-d} < \bar{\lambda}$ and $1-\bar{\lambda} < \frac{c-a}{c-a+b-d} < 1-\underline{\lambda}$ are not the case and suppose, towards a contradiction, that the policy σ^* such that $\sigma^*(Mm) = I$ and $\sigma^*(Reg) = II$ defines a symmetric Γ -equilibrium $\vec{\sigma}^* = (\sigma^*, \sigma^*)$. For that to be the case, it must hold that, for $i \in \{1, 2\}$:

$$\min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((I, II), \sigma^*))|Mm] \geq \min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((II, II), \sigma^*))|Mm], \quad (6.9)$$

and that

$$\min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((I, II), \sigma^*))|Reg] \geq \min_{\lambda} \mathbb{E}_{\lambda}[u_i(\pi((I, I), \sigma^*))|Reg] \quad (6.10)$$

Given $\lambda \in \Gamma = [\underline{\lambda}, \bar{\lambda}] \subseteq [0, 1]$, the expected fitness of action I associated to $\vec{\sigma}^*$ is:

$$\mathbb{E}_{\lambda}[\Phi(\pi(I, \sigma^*))] = (1 - \lambda)a + \lambda b,$$

while the expected fitness of action II is:

$$\mathbb{E}_{\lambda}[\Phi(\pi(II, \sigma^*))] = (1 - \lambda)c + \lambda d.$$

Consequently, action I has higher expected fitness than action II if

$$(1 - \lambda)a + \lambda b - (1 - \lambda)c - \lambda d > 0$$

which is equivalent to $\lambda < \frac{c-a}{c-a+b-d}$. In the same way, action II has higher expected fitness than action I if $\lambda > \frac{c-a}{c-a+b-d}$. Remember that, for any single λ ,

$$\mathbb{E}_{\lambda}[\Phi(\pi(I, \sigma^*))] \geq \mathbb{E}_{\lambda}[\Phi(\pi(II, \sigma^*))]$$

iff

$$\mathbb{E}_{\lambda}[\Phi(\pi(I, \sigma^*))] - \max_{a' \in A} \mathbb{E}_{\lambda}[\Phi(\pi(a', \sigma^*))] \geq \mathbb{E}_{\lambda}[\Phi(\pi(II, \sigma^*))] - \max_{a' \in A} \mathbb{E}_{\lambda}[\Phi(\pi(a', \sigma^*))].$$

This is, again, the reason why maximinimizers and regret minimizers do not differ when Γ is a singleton set. But then if $\frac{c-a}{c-a+b-d} < \underline{\lambda}$, action II is strongly dominated by action I for all $\lambda \in \Gamma$ for both types, which contradicts equation (6.10), so

that type *Reg* will have incentives to deviate from the policy σ^* . Similarly, if $\bar{\lambda} < \frac{c-a}{c-a+b-d}$ then action *I* is strongly dominated by action *II* for all $\lambda \in \Gamma$ for both types, which contradicts equation (6.9), so that type *Mm* will have incentives to deviate from the policy σ^* . Hence, the revealing profile $\vec{\sigma}^*$ such that $\sigma^*(Mm) = I$ and $\sigma^*(Reg) = II$ cannot be a Γ -equilibrium of the game. Analogously for the other revealing profile $\vec{\sigma}^* = (\sigma^*, \sigma^*)$ such that $\sigma^*(Mm) = II$ and $\sigma^*(Reg) = I$, if $1 - \bar{\lambda} < \frac{c-a}{c-a+b-d} < 1 - \underline{\lambda}$ is not the case. \square

The following is the proof of Lemma 6.5.

Proof. As shown in Figure 6.2a, in an anti-coordination game action *I* corresponds to the line $a + (b-a)x$, while action *II* corresponds to the line $c + (d-c)x$. The slope of action *I* is then $(b-a)$, and the slope of action *II* is $(d-c)$. Action *I* is steeper than action *II* if $|b-a| > |d-c|$, and action *II* is steeper than action *I* if the reverse of the last inequality holds. Given a belief $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ and a policy function σ , let us define the following preference relations over actions. The relation $\succsim_{M|\Gamma, \sigma}$ corresponding to type *Mm* is defined by

$$I \succsim_{M|\Gamma, \sigma} II \text{ iff } \min_{\lambda} \mathbb{E}_{\lambda}[\Phi(\pi(I, \sigma))] \geq \min_{\lambda} \mathbb{E}_{\lambda}[\Phi(\pi(II, \sigma))],$$

while the relation $\succsim_{R|\Gamma, \sigma}$ corresponding to type *Reg* is defined by

$$\begin{aligned} I \succsim_{R|\Gamma, \sigma} II \text{ iff} \\ \min_{\lambda} \{ \mathbb{E}_{\lambda}[\Phi(\pi(I, \sigma))] - \max_{a' \in A} \mathbb{E}_{\lambda}[\Phi(\pi(a', \sigma))] \} \\ \geq \\ \min_{\lambda} \{ \mathbb{E}_{\lambda}[\Phi(\pi(II, \sigma))] - \max_{a' \in A} \mathbb{E}_{\lambda}[\Phi(\pi(a', \sigma))] \}. \end{aligned}$$

For any fixed belief Γ and policy σ , it is the case that $|d-c| > |b-a|$ entails

$$I \succsim_{R|\Gamma, \sigma} II \Rightarrow I \succ_{M|\Gamma, \sigma} II.$$

Indeed, when a belief $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ about types' proportions is paired with a policy σ , the two induce a belief $\Gamma|\sigma$ over the co-player's actions in the obvious way. For a given $\Gamma|\sigma = [\underline{p}, \bar{p}]$, where \underline{p} is now the lower probability of action *II* and \bar{p} is the upper probability of action *II*, let us define

$$\begin{aligned} a' &:= (1 - \underline{p})a + \underline{p}b = a + \underline{p}(b-a) \\ b' &:= (1 - \bar{p})a + \bar{p}b = a + \bar{p}(b-a) \\ c' &:= (1 - \underline{p})c + \underline{p}d = c + \underline{p}(d-c) \\ d' &:= (1 - \bar{p})c + \bar{p}d = c + \bar{p}(d-c) \end{aligned}$$

Next, type *Reg* is indifferent between the two acts, $I \sim_{R|\Gamma, \sigma} II$, if $\frac{c-a}{c-a+b-d} - \underline{p} = \bar{p} - \frac{c-a}{c-a+b-d}$, and prefers *I* over *II* if $\frac{c-a}{c-a+b-d} - \underline{p} < \bar{p} - \frac{c-a}{c-a+b-d}$. For succinctness, let

us abbreviate $Z := c - a + b - d$. Whenever $|d - c| > |b - a|$ and $\frac{c-a}{Z} - \underline{p} \leq \bar{p} - \frac{c-a}{Z}$, it is the case that $a' > d'$, so that $I \succ_{M|\Gamma, \sigma} II$. Indeed, when $\underline{p} = \bar{p} = \frac{c-a}{Z}$, we have that $a' = d' = \frac{cb-ad}{Z}$. When we enlarge the interval $\Gamma|\sigma$ by moving \underline{p} to the left of $\frac{c-a}{Z}$ and \bar{p} to the right of $\frac{c-a}{Z}$ by the same extent, such that $\frac{c-a}{Z} - \underline{p} = \bar{p} - \frac{c-a}{Z}$, we get that $a' > d'$, since a' moved by $(a-b)(\frac{c-a}{Z} - \underline{p})$ while d' moved by $(d-c)(\bar{p} - \frac{c-a}{Z})$. Consequently, in an anti-coordination game where action II is steeper than action I , the only possible revealing symmetric Γ -equilibrium is a profile $\bar{\sigma}^* = (\sigma^*, \sigma^*)$ such that $\sigma^*(Mm) = I$ and $\sigma^*(Reg) = II$. By a similar argument, when action I is steeper than action II , the only possible revealing symmetric Γ -equilibrium is a profile $\bar{\sigma}^* = (\sigma^*, \sigma^*)$ such that $\sigma^*(Mm) = II$ and $\sigma^*(Reg) = I$. Hence, whenever one of the actions is steeper than the other, only one of the two revealing policies can constitute a Γ -equilibrium, so that we excluded the case of games with two revealing symmetric Γ -equilibria.

We are left with the case where action I and action II are equally steep, $|d - c| = |b - a|$. We will see later that these are games such that $\frac{b-a}{Z} = \frac{1}{2}$, and they are never strongly informative. Indeed, by the same argument as above we obtain that $|d - c| = |b - a|$ implies

$$I \succ_{R|\Gamma, \sigma} II \Leftrightarrow I \succ_{M|\Gamma, \sigma} II,$$

so that, for each Γ and σ , the two types will always choose the same action. \square

Finally, we can prove Proposition 6.6.

Proof. Only if. Notice first that, by simple computation, the condition expressed by inequality 6.2 is equivalent to the condition

$$\frac{c-a}{Z} < \frac{\bar{\lambda}(c-d) + \underline{\lambda}(b-a)}{Z}.$$

This allows us to formulate both condition 6.2 and condition 6.3 in terms of the intersection point $\frac{c-a}{Z}$ between line I and line II , and to put them together in a unique double inequality, as done in the first point of Proposition 6.6:

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c-a}{Z} < \frac{\bar{\lambda}(c-d) + \underline{\lambda}(b-a)}{Z}.$$

In the same way, the condition expressed by inequality 6.5 is equivalent to

$$\frac{(1-\underline{\lambda})(c-d) + (1-\bar{\lambda})(b-a)}{Z} < \frac{c-a}{Z},$$

so that we can bring condition 6.5 and condition 6.6 together:

$$\frac{(1-\underline{\lambda})(c-d) + (1-\bar{\lambda})(b-a)}{Z} < \frac{c-a}{Z} < \frac{(1-\bar{\lambda}) + (1-\underline{\lambda})}{2},$$

as in the second point of Proposition 6.6.

Now, by contraposition, suppose that both conditions of Proposition 6.6 are violated, that is, there is some interval $[\underline{\gamma}, \bar{\gamma}] \subseteq [0, 1]$ (with $\underline{\gamma} < \bar{\gamma}$) such that there are no anti-coordination games (a, b, c, d) for which one of the two conditions in Proposition 6.6 holds. Then, whenever the interval $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ hits the interval $[\underline{\gamma}, \bar{\gamma}]$, the players will no longer observe any informative plays in the future. In particular, if the first condition of Proposition 6.6 does not hold for some interval $[\underline{\gamma}, \bar{\gamma}] \subseteq [0, 1]$, it follows that there will be no more anti-coordination games where the policy (I, II) constitutes a symmetric Γ -equilibrium for $\Gamma \subseteq [\underline{\gamma}, \bar{\gamma}]$. Reversely, if the second condition does not hold either, there will be no anti-coordination games where the dual policy (II, I) constitutes a symmetric Γ -equilibrium for $\Gamma \subseteq [\underline{\gamma}, \bar{\gamma}]$. Suppose now that the true distribution λ^* of regret types lies in that interval, $\underline{\gamma} < \lambda^* < \bar{\gamma}$. From the assumption that both conditions of Proposition 6.6 are violated, it follows that, whenever $\Gamma \subseteq [\underline{\gamma}, \bar{\gamma}]$, both types will play the same action in all future games. Consequently, the players will no longer face any strongly informative game in the future, and they will not be able to shrink the set Γ any further, so that they are prevented from learning the true distribution λ^* .

If. By the same reasoning, if for any interval $[\underline{\lambda}, \bar{\lambda}] \subseteq [0, 1]$ at least one of the two conditions of Proposition 6.6 is satisfied, then it means that whatever the belief interval $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ is, the players will always face, sooner or later, a strongly informative game, i.e., a game that satisfies one of the conditions of Proposition 6.6. Specifically, when the first condition is satisfied, this ensures that there will be anti-coordination games for which the policy function (I, II) constitutes the unique symmetric Γ -equilibrium; when the second condition is fulfilled, there will be games where the dual policy (II, I) defines the unique symmetric Γ -equilibrium. In such games, the two types would play two different actions, so that the players will be able to further shrink the interval Γ , according to the behavior observed in those games. If strongly informative games will always arise, independent of the current interval Γ , then the interval Γ can always be shrunk further and further, and the players will eventually be able to learn the true distribution λ^* in the limit. \square

The following is the proof of Corollary 6.11.

Proof. Proposition 6.6 gives necessary and sufficient conditions for learnability in anti-coordination games. We now want to show that if games come from i.i.d. drawings of a continuous random variable, then each of the two conditions of Proposition 6.6 follows. (By similar arguments it is possible to show that the same holds for the two conditions of Proposition 6.10 for coordination games.) Suppose that games (a, b, c, d) are generated by i.i.d drawings of a continuous random variable with probability distribution P_V , and that the real interval (\underline{r}, \bar{r}) is a subset of its support. Fix any set $\Gamma = [\underline{\lambda}, \bar{\lambda}]$, corresponding to the uncertainty

of the players about the distribution of types, and consider the first condition of Proposition 6.6:

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c - a}{Z} < \frac{\bar{\lambda}(c - d) + \underline{\lambda}(b - a)}{Z}. \quad (6.11)$$

Note also that the ratio $\frac{c-a}{Z}$ takes values in the interval $(0, 1)$, since to get $\frac{c-a}{Z} = 0$ we shall have $a = c$ and for $\frac{c-a}{Z} = 1$ we need $d = b$, which have no probability of occurrence for continuous random variables. Suppose that two random numbers $c, d \in (\underline{r}, \bar{r})$ are drawn. Two cases are equiprobable: $c > d$ and $c < d$. Consider the first, $c > d$. Since $(d, c) \subseteq (\underline{r}, \bar{r})$, the event that $a \in (d, c)$ has positive probability, i.e., there is positive probability of drawing a such that $d < a < c$.

Take the interval $\left[\frac{\underline{\lambda} + \bar{\lambda}}{2}, \bar{\lambda}\right]$, and consider the two extrema. The line passing through a and intersecting the line corresponding to action II passing through c and d at $x = \frac{\underline{\lambda} + \bar{\lambda}}{2}$ has equation:

$$y = x \left(\frac{c + \frac{\underline{\lambda} + \bar{\lambda}}{2}(d - c) - a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} \right) + a,$$

while the line passing through a and intersecting the line corresponding to action II passing through c and d at $x = \bar{\lambda}$ has equation:

$$y = x \left(\frac{c + \bar{\lambda}(d - c) - a}{\bar{\lambda}} \right) + a.$$

If we consider the values of y for $x = 0$ and $x = 1$, each of these lines corresponds to possible fitness numbers for action I (each line gives a for $x = 0$ and b for $x = 1$).

Each point $z \in \left(\frac{\underline{\lambda} + \bar{\lambda}}{2}, \bar{\lambda}\right)$ can be written as a convex combination

$$z = \alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1 - \alpha)\bar{\lambda}$$

for some $\alpha \in (0, 1)$. So, the line passing through a and intersecting the line corresponding to action II at $x = \alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1 - \alpha)\bar{\lambda}$ has equation:

$$y = x \left(\frac{c + \left(\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1 - \alpha)\bar{\lambda}\right)(d - c) - a}{\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1 - \alpha)\bar{\lambda}} \right) + a.$$

As we have seen, each of these lines corresponds to a value of b for $x = 1$. The set of all these values of b is thus the set

$$B := \left\{ \frac{c - a}{\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1 - \alpha)\bar{\lambda}} + d - c + a : \alpha \in (0, 1) \right\}.$$

Since all these values are included in the interval

$$\left(\frac{c-a}{\bar{\lambda}} + d - c + a, \frac{c-a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} + d - c + a \right),$$

then we can express the set B as a set of convex combinations between these two extreme points:

$$(1-\alpha) \left[\frac{c-a}{\bar{\lambda}} + d - c + a \right] + \alpha \left[\frac{c-a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} + d - c + a \right].$$

Notice, however, that in general

$$(1-\alpha) \left[\frac{c-a}{\bar{\lambda}} + d - c + a \right] + \alpha \left[\frac{c-a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} + d - c + a \right] \neq \frac{c-a}{\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1-\alpha)\bar{\lambda}} + d - c + a,$$

and, more precisely, we have that

$$\begin{aligned} (1-\alpha) \left[\frac{c-a}{\bar{\lambda}} + d - c + a \right] + \alpha \left[\frac{c-a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} + d - c + a \right] & - \frac{c-a}{\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1-\alpha)\bar{\lambda}} + d - c + a \\ & = \\ & (c-a) \frac{\alpha(1-\alpha) \left(\bar{\lambda} - \frac{\underline{\lambda} + \bar{\lambda}}{2} \right)^2}{\underline{\lambda} \bar{\lambda} \left(\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1-\alpha)\bar{\lambda} \right)}. \end{aligned}$$

Hence, we can see that in general the difference between the two is always positive (since $c > a$ by assumption, and $\bar{\lambda} > \frac{\underline{\lambda} + \bar{\lambda}}{2}$), so that

$$(1-\alpha) \left[\frac{c-a}{\bar{\lambda}} + d - c + a \right] + \alpha \left[\frac{c-a}{\frac{\underline{\lambda} + \bar{\lambda}}{2}} + d - c + a \right] > \frac{c-a}{\alpha \frac{\underline{\lambda} + \bar{\lambda}}{2} + (1-\alpha)\bar{\lambda}} + d - c + a.$$

Moreover, as one would expect, the difference between the two tends to 0 both when α tends to 0 and when α tends to 1.

This shows that the set B corresponds to an interval $B \subset (d, c) \subset (\underline{r}, \bar{r})$. Consequently, for any $(d, c) \subseteq (\underline{r}, \bar{r})$, there is always positive probability to find $a, b \in (d, c)$ such that the game (a, b, c, d) is strongly informative. \square

*Life is not a matter of holding good cards,
but of playing a poor hand well.* (R. L. Stevenson)

7.1 The Game of Life with Learning

The obvious question that the reader should ask himself or herself at this point is: what happens when we consider a population that plays the game of life, where players *both* learn according to the learning dynamics described in Chapter 6 *and* evolve according to the evolutionary dynamics of Chapter 5? This is a relevant and legitimate question: it is possible that the stability of some evolutionary outcomes can be destabilized by the introduction of learning.

In the new game of life, enriched with the possibility of learning, players are involved in two different dynamics. We think of the situation as follows. During their life, within each generation, players face many different interactive decision problems. The intragenerational dynamics are exactly as described by points 1-8 of Section 6.1: at each time t a game is selected according to some probability P_G over all possible games, two players from the population are randomly matched, the two players play the game, and the outcome is commonly known within the population. Moreover, in order to add the intergenerational evolutionary dynamics to the process, we need to keep track of the accumulated fitness of different types, as we did in Chapter 5. As $t \rightarrow \infty$, players might eventually be able to learn the precise distribution of types in the population (depending on the games in the environment, as we have seen in Chapter 6), and will also accumulate some evolutionary fitness (depending on their type, as we have seen in Chapter 5). In particular, at any moment $t' < \infty$, all types in the population will have an associated accumulated fitness on the basis of their (expected) fitness

in all past $t' - 1$ games. Natural selection is then based on the fitness accumulated by each type, as in Chapter 5.

7.2 Play Without Regret

Let us take into consideration a population with the two decision types Mm and Reg , i.e., a population where players can only differ in the decision criterion \hat{a} . More specifically, the subjective utility u^j of all types $t^j = (\hat{a}^j, u^j, \Gamma^j) \in \mathcal{T}$ is anchored to objective fitness, $u^j = \Phi$ for all $t^j \in \mathcal{T}$, and, at any time t , $\Gamma^j = \Gamma^t$ for all $t^j \in \mathcal{T}$, that is: all players hold the same confidence interval, in line with the learning dynamics of Chapter 6. Consequently, we can identify different phenotypes in the population by their decision criterion, i.e., $\mathcal{T} = \{Mm, Reg\}$. Possible differences in action choices will only depend on different decision criteria. *Ceteris paribus*, we are then looking at the evolution of decision criteria here, leaving aside possible variability in subjective utilities and beliefs for the moment. The departure from Chapter 5 consists in that the set Γ^j is now subject to a process of repeated updates.

From these premises, it is possible to proceed to the evolutionary analysis of the game of life with learning. Intragenerational dynamics are as described in points 1-8 of Section 6.1. The intergenerational dynamics are given by the evolutionary dynamics:

9. At the end of each generation at time $T < \infty$, each type accumulated a certain (expected) fitness, which determines the reproduction of the type in the next generation.
10. After time T , a new generation of players is created, and begins to play the game of life G^1, G^2, G^3, \dots from $t = 0$ again, according to points 1-8.

We refer to the evolutionary multi-game specified by points 1-10 as the *game of life with learning*.

The following proposition extends the evolutionary results of Chapter 5 to the game of life with learning.

Proposition 7.1. *In the game of life with learning generated by drawing a continuous random variable with uniform distribution on a given support, Reg strictly dominates Mm .*

Proof. In Appendix 7.5. □

The following corollary confirms the evolutionary results of Corollary 5.5 also for the case of the game of life with learning.

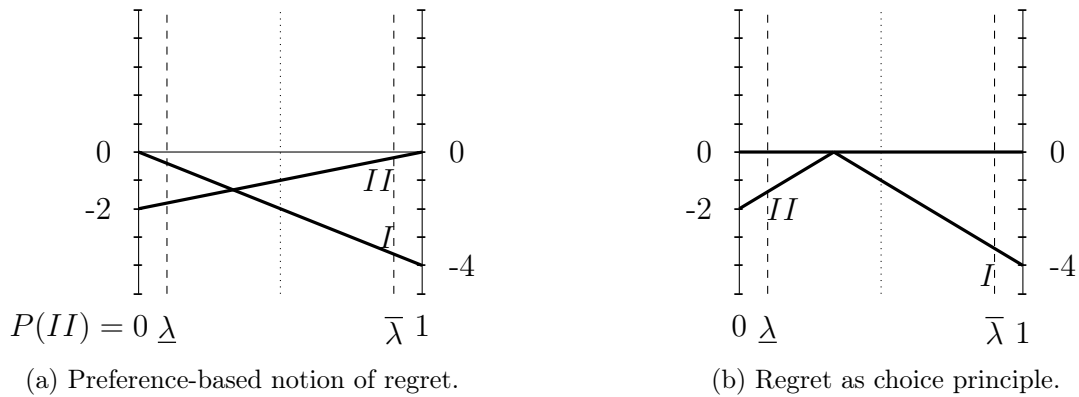


Figure 7.1: The two possible representations of regret.

Corollary 7.2. *Fix a population of phenotypes Mm and Reg . In the game of life with learning generated by drawing a continuous random variable with uniform distribution on a given support, the only evolutionarily stable state is a monomorphic population of type Reg .*

7.3 Regret What?

In Section 5.9.1, we showed how different pairings of decision criteria and subjective utilities may generate the same indistinguishable behavior. In particular, we showed that it would have been possible to represent simple regret minimization both by using regret-based preferences coupled with maxmin decision criterion, and by regret minimization acting on objective preferences. The resulting behavior would have been the same. This is no longer the case once we introduce learning, and uncertainty is not necessarily maximal.

Consider the population game of Section 3.3. There, we opted for the preference-based representation of regret: type t^2 associates each outcome with the regret at that outcome, i.e., $u_{t^2}(\pi(a, a')) = \text{reg}(\pi(a, a'))$. On the other hand, in the last two chapters, we expressed regret directly as a decision criterion rather than a subjective preference.

For the case of the population game in Section 3.3, Figure 7.1 graphically represents the difference between formalizing regret as a decision criterion and formalizing regret in terms of a subjective utility function. The picture on the right corresponds to the subjective preference version of regret, defined as the difference of single outcomes:

$$Mm(\text{reg}, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \mathbb{E}_P[u(\pi(a, a')) - \max_{a'' \in A} u(\pi(a'', a'))].$$

The picture on the left instead represents regret as a decision criterion, as defined

in Chapter 5:

$$Reg(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \{ \mathbb{E}_P[u(\pi(a, a'))] - \max_{a'' \in A} \mathbb{E}_P[u(\pi(a'', a'))] \}.$$

In the first case, we say that regret is expressed as a subjective preference because it is possible to express it through a subjective utility $\operatorname{reg} : X \rightarrow \mathbb{R}$, whereas this is not possible in the second case.

Notice from Figure 7.1 that both notions of regret would dictate action II : the minimal expected utility of I in the interval $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ is lower than the minimal expected utility of II both in 7.1a and 7.1b. This is not true in general, though.

Suppose there is a possible bet on an urn which contains ten balls that are either white or black. The bet pays \$1000 if a white ball is drawn and the bettor guesses the color, \$3000 if a black ball is drawn and the bettor guesses the color, and nothing otherwise, as shown in the table below.

	W	B
w	1000	0
b	0	3000

It is known that the urn contains one black ball and seven white balls, but the color of the last two balls is unknown (this is the case of $\Gamma = [\underline{\lambda}', \bar{\lambda}]$ in Figure 7.2). Preference-based regret is indifferent between betting on white and betting on black in such a situation, while the second account of regret would prefer to bet on white.

If, instead, it is known that two balls are black, seven are white, and the color of only one ball is unknown, then preference-based regret would bet on black, while the other regret would be indifferent between the two colors (this is the case of $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ in Figure 7.2).

Figure 7.2 represents the two different situations, in terms of preference-based regret (on the left), and in terms of regret as decision criterion (on the right). Clearly, the regret representation on the right is indifferent between the two actions whenever the set of parameters $\Gamma = [\underline{\lambda}, \bar{\lambda}]$ is symmetric around $\frac{1}{4}$: this is the case of the urn where the color of only one ball is unknown. When the set Γ is not symmetric around $\frac{1}{4}$, instead, regret-choice principle strictly prefers one of the two actions. Hence, when preference-based regret is indifferent between b and w , i.e., when $\underline{\lambda} = 0.1$ and $\bar{\lambda} = 0.3$, regret-choice principle will prefer action w over b . Reversely, when regret-choice principle is indifferent, preference-based regret will uniquely choose action b .

One final remark is in order here. Notice that the axiomatic system given by [Stoye, 2011] represents the preference-based notion of regret, which is not the choice principle version of regret that we used for our results about evolution and learning in the last chapters. Opting for one notion of regret over the other, rather than necessary, is a matter of preference and intuition, since both preference-based




Figure 7.2: The two possible representations of regret.

regret and regret-choice principle express reasonable and viable norms for decision making. It is left to experimental work to determine if agents tend to resort to one more than the other.

7.4 Conclusion

The main goal of this chapter was to combine the results about the evolutionary competition from Chapter 5 with the results about the learning from Chapter 6. This unified perspective allows us to consider the evolution of more sophisticated agents, that are able to learn about the proportions of decision types within their generation while playing the game of life. The main conclusion that the presence of unmeasurable uncertainty generally favors regret minimizing players over maximinimizers hasn't been destabilized by the introduction of the learning dynamics.

7.5 Appendix: Proofs

Definition 7.3. φ . The permutation $\varphi(a, b, c, d) = (d, c, b, a)$ defines a bijective function from coordination games to coordination games $\varphi : \mathcal{C} \rightarrow \mathcal{C}$, and from anti-coordination games to anti-coordination games $\varphi : \mathcal{A} \rightarrow \mathcal{A}$.

Lemma 7.4. For coordination and anti-coordination games $G \in \mathcal{C} \cup \mathcal{A}$, the policy (I, II) constitutes a symmetric Γ -equilibrium in game G if and only if the dual policy (II, I) constitutes a symmetric Γ -equilibrium in game $\varphi(G)$.

Proof. Geometrically, the bijection φ amounts to a reflection about the y -axis, plus a relabeling of the actions. Suppose now that the policy (I, II) defines a

(strict) symmetric Γ -equilibrium for the anti-coordination game A . For that to be the case, we know from Proposition 6.6 that it must hold that

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c - a}{Z} < \frac{\bar{\lambda}(c - d) + \underline{\lambda}(b - a)}{Z},$$

so that I is a best reply for type Mm , and II is a best reply for type Reg . In the corresponding game $\varphi(A)$, (by simple relabeling) it must hold that:

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{b - d}{Z} < \frac{\bar{\lambda}(b - a) + \underline{\lambda}(c - d)}{Z}.$$

But notice that

$$\frac{b - d}{Z} = 1 - \frac{c - a}{Z},$$

and hence we can rewrite the previous double inequality as

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < 1 - \frac{c - a}{Z} < \frac{\bar{\lambda}(b - a) + \underline{\lambda}(c - d)}{Z}.$$

By computation, the last reduces to

$$\frac{(1 - \bar{\lambda})(b - a) + (1 - \underline{\lambda})(c - d)}{Z} < \frac{c - a}{Z} < \frac{(1 - \underline{\lambda}) + (1 - \bar{\lambda})}{2},$$

which corresponds to the necessary and sufficient conditions for $((II, I), (II, I))$ to be a Γ -equilibrium of the anti-coordination game $\varphi(A)$. Hence, the policy (I, II) defines a symmetric Γ -equilibrium in game A if and only if the dual policy (II, I) defines a symmetric Γ -equilibrium in game $\varphi(A)$. By analogous argument, the same holds for coordination games too. \square

We can finally prove Proposition 7.1.

Proof. First of all, notice that the revealing Γ -equilibria $((I, II), (I, II))$, where Mm plays I and Reg plays II , are only possible for $\underline{\lambda} < \frac{c-a}{Z} < \bar{\lambda}$, since λ is the proportion of regret types in the population. Reversely, revealing equilibria $((II, I), (II, I))$ are only possible if $1 - \bar{\lambda} < \frac{c-a}{Z} < 1 - \underline{\lambda}$. Then, by Lemma 7.4 we can reason only in terms of the interval $[\underline{\lambda}, \bar{\lambda}]$ and of Γ -equilibria of the form $((I, II), (I, II))$, knowing that for any (anti-)coordination game that has a revealing Γ -equilibrium $((I, II), (I, II))$ there is a corresponding (anti-)coordination game such that $1 - \bar{\lambda} < \frac{c-a}{Z} < 1 - \underline{\lambda}$ and such that $((II, I), (II, I))$ is a Γ -equilibrium where players get exactly the same fitness as in the Γ -equilibrium $((I, II), (I, II))$ of the original game. By this move we can reduce the analysis just to the interval $[\underline{\lambda}, \bar{\lambda}]$ and to Γ -equilibria of the form $((I, II), (I, II))$.




Figure 7.3: Examples of revealing games with equilibrium $((I, II), (I, II))$.

Second, by the conditions of Proposition 6.6 and Proposition 6.10 and by Lemma 6.4, it follows that, for any given interval $[\underline{\lambda}, \bar{\lambda}]$, the revealing coordination games with Γ -equilibrium $((I, II), (I, II))$ must be such that

$$\underline{\lambda} < \frac{c - a}{Z} < \frac{\underline{\lambda} + \bar{\lambda}}{2},$$

whereas the revealing anti-coordination games with Γ -equilibrium $((I, II), (I, II))$ must be such that

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c - a}{Z} < \bar{\lambda}.$$

Figure 7.3 shows examples of revealing coordination and anti-coordination games for a given interval $[\underline{\lambda}, \bar{\lambda}]$. As we can see, the two lines intersect on the left of the middle of the interval in the coordination game, and on the right in the anti-coordination game.

From Lemma 6.7 we know that a necessary condition for $((I, II), (I, II))$ to be a Γ -equilibrium is that

$$|d - c| > |a - b|.$$

Suppose now that the two numbers c and d have been drawn. From the last inequality we already know that if $c < d$ then the only possible revealing games in $[\underline{\lambda}, \bar{\lambda}]$ are coordination games, because an anti-coordination games by definition satisfies $a < c$ and $d < b$, which contradicts the necessary condition $|d - c| > |a - b|$. Reversely, if $c > d$, then the only revealing games in $[\underline{\lambda}, \bar{\lambda}]$ are anti-coordination games. Without loss of generality, suppose that c and d are drawn such that $c < d$, so that the only possible revealing games in $[\underline{\lambda}, \bar{\lambda}]$ are coordination games. By the argument above, the revealing coordination games in $[\underline{\lambda}, \bar{\lambda}]$ must be such

that

$$\underline{\lambda} < \frac{c - a}{Z} < \frac{\underline{\lambda} + \bar{\lambda}}{2},$$

since this is the necessary and sufficient condition for *Reg* to (strictly) prefer *II* over *I*. Consider now all the points $\lambda \in (\underline{\lambda}, \frac{\underline{\lambda} + \bar{\lambda}}{2})$. Each of these points can be expressed as a linear combination:

$$\frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n},$$

for $k > \frac{n}{2}$. By simple algebra, for each point in $\lambda \in (\underline{\lambda}, \frac{\underline{\lambda} + \bar{\lambda}}{2})$ and $c < d$, the point

$$c + \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c)$$

expresses the expected value of action *II* for $\lambda = \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}$, i.e., it is the *y*-value of the line corresponding to action *II* when $x = \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}$. Then, given a two-dimensional point (x_0, y_0) , the sheaf of lines passing through that point is defined by all the equations

$$y - y_0 = m(x - x_0)$$

for $m \in \mathbb{R}$. (The vertical line $m = \infty$ is excluded from the sheaf, but it is not relevant for the proof.) Consequently, given the two-dimensional point

$$\left(\frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}, c + \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) \right),$$

the sheaf of lines passing through that point is defined by the set of equations, for $m \in \mathbb{R}$:

$$y - c - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) = m \left(x - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right).$$

If, for each equation in the set, we define

$$\begin{aligned} a^\diamond &:= m \left(-\frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right) + c + \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) \\ b^\diamond &:= m \left(1 - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right) + c + \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) \end{aligned} \quad (7.1)$$

then each equation corresponds to a possible game such that

$$\frac{c - a^\diamond}{c - a^\diamond + b^\diamond - d} = \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}.$$

Among those, we are interested in all the games that are revealing, i.e., the games for which $((I, II), (I, II))$ is a Γ -equilibrium. From the previous analysis, we

know that only coordination games such that $d - c > |a^\diamond - b^\diamond|$ can be informative. Indeed, if $a^\diamond < b^\diamond$ and $d - c < b^\diamond - a^\diamond$, then the game is an anti-coordination game, and it cannot possibly support the profile $((I, II), (I, II))$ as a Γ -equilibrium for $\frac{c - a^\diamond}{c - a^\diamond + b^\diamond - d} \in (\underline{\lambda}, \frac{\underline{\lambda} + \bar{\lambda}}{2})$; if instead $a^\diamond > b^\diamond$ and $d - c < a^\diamond - b^\diamond$, then the game is still a coordination game, but we have that $|d - c| < |a^\diamond - b^\diamond|$, and according to Lemma 6.7 the profile $((I, II), (I, II))$ cannot be a Γ -equilibrium of the game. By algebraic computations, the condition $d - c > |a^\diamond - b^\diamond|$ is equivalent to

$$|m| < d - c.$$

Moreover, among the coordination games such that $d - c > |a^\diamond - b^\diamond|$, the informative ones are those that also satisfy $b' > c'$, otherwise type Mm would not (strictly) prefer I over II . If we rewrite a', b', c', d' as in equation 6.1, then the inequality $b' > c'$ reduces to

$$m \left(\frac{k}{k - n} \right) < d - c.$$

To sum up, so far we have seen that each line

$$y - c - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) = m \left(x - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right)$$

corresponds to a possible game:

	I	II
I	a^\diamond	b^\diamond
II	c	d

such that

$$\frac{c - a^\diamond}{c - a^\diamond + b^\diamond - d} = \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}.$$

When $d > c$, the only possible revealing games are coordination games, since it is necessary $|d - c| > |a^\diamond - b^\diamond|$ for the profile $((I, II), (I, II))$ to be an equilibrium. Furthermore, for a coordination game to be revealing it must also hold that $b' > c'$. Hence, in the end, we are looking for all games

$$y - c - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) = m \left(x - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right)$$

with the constraints that

$$|m| < d - c$$

and

$$m \left(\frac{k}{k - n} \right) < d - c.$$

By symmetric arguments, whenever $c > d$, the only possible revealing games for the same interval $[\underline{\lambda}, \bar{\lambda}]$ are anti-coordination games for which

$$\frac{\underline{\lambda} + \bar{\lambda}}{2} < \frac{c - a}{Z} < \bar{\lambda},$$

and such that $c - d > |a - b|$, and that $a' > d'$. Similarly, for $k < \frac{n}{2}$, these correspond to all games

$$y - c - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) = m \left(x - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right)$$

such that

$$|m| < c - d$$

and

$$m \left(\frac{k - n}{k} \right) > d - c.$$

Consider now the following bijective function $\psi : \mathcal{C} \rightarrow \mathcal{A}$ between coordination and anti-coordination games, that, for $d > c$, associates the coordination game

$$y - c - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n}(d - c) = m \left(x - \frac{k\underline{\lambda} + (n - k)\bar{\lambda}}{n} \right)$$

with the anti-coordination game

$$y - d - \frac{(n - k)\underline{\lambda} + k\bar{\lambda}}{n}(c - d) = -m \left(x - \frac{(n - k)\underline{\lambda} + k\bar{\lambda}}{n} \right).$$

Essentially, ψ changes c to d , m to $-m$, and k to $n - k$. In particular, note that ψ is a bijection that, for a fixed interval $[\underline{\lambda}, \bar{\lambda}]$, sends revealing coordination games to revealing anti-coordination games. Figure 7.4 gives a graphical example of the bijection.

We can then pair these two games and consider the average fitness in $\{C, \psi(C)\}$ of a type *Reg* against another type *Reg*, and then compare it to the fitness of a type *Mm* against a type *Reg*. In the pair of revealing games C and $\psi(C)$, *Reg* strictly dominates *Mm* if

$$F_{\{C, \psi(C)\}}(R, R) > F_{\{C, \psi(C)\}}(M, R)$$

and

$$F_{\{C, \psi(C)\}}(R, M) > F_{\{C, \psi(C)\}}(M, M).$$

Consider the first inequality. Since both C and $\psi(C)$ are supposed to be revealing with respect to the interval $[\underline{\lambda}, \bar{\lambda}]$, it implies that $F_{\{C, \psi(C)\}}(R, R) = d + c$, and $F_{\{C, \psi(C)\}}(M, R) = b^\diamond + \psi(b^\diamond)$. Therefore, the first inequality is equivalent to

$$\frac{d + c}{2} > \frac{b^\diamond + \psi(b^\diamond)}{2},$$

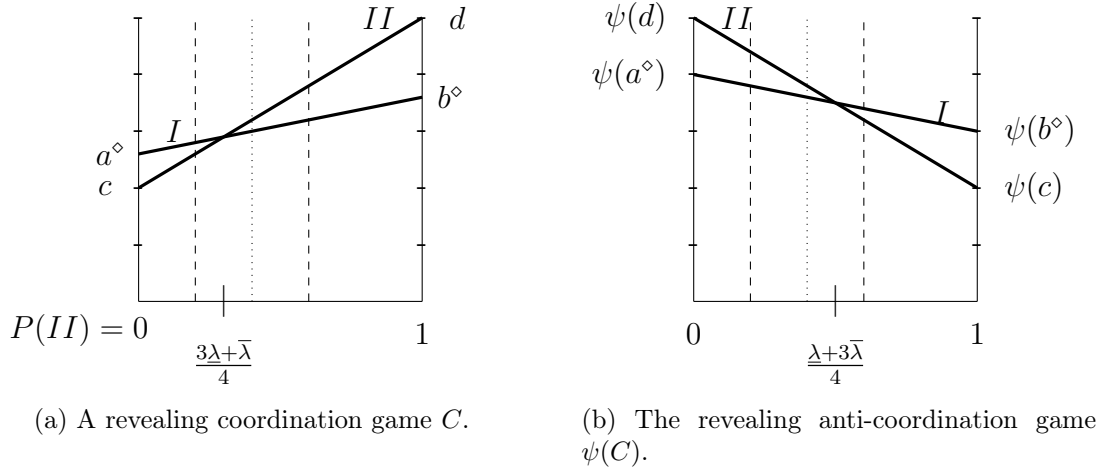


Figure 7.4: Examples of revealing games with equilibrium $((I, II), (I, II))$ under the bijection ψ .

which can be spelled out as

$$\begin{aligned} d + c &> m \left(1 - \frac{k\lambda + (n-k)\bar{\lambda}}{n} \right) + c + \frac{k\lambda + (n-k)\bar{\lambda}}{n} (d - c) \\ &\quad - m \left(1 - \frac{(n-k)\lambda + k\bar{\lambda}}{n} \right) + d + \frac{(n-k)\lambda + k\bar{\lambda}}{n} (c - d) \end{aligned}$$

After some computations, the previous inequality boils down to

$$d - c > m,$$

which we know it is the case, since we have seen that the condition $d - c > |a^\diamond - b^\diamond|$ is equivalent to $d - c > |m|$.

Finally, from the previous argument it follows that, for any given interval $[\underline{\lambda}, \bar{\lambda}]$, if we consider the set of all revealing coordination games, let's denote it by \mathcal{C}^r , and the set of all revealing anti-coordination games \mathcal{A}^r , then it holds that

$$F_{\{\mathcal{C}^r \cup \mathcal{A}^r\}}(R, R) > F_{\{\mathcal{C}^r \cup \mathcal{A}^r\}}(M, R).$$

Let us now check that the second inequality for Reg to strictly dominate Mm also holds. In $\{C, \psi(C)\}$, that is equivalent to

$$d + c > a^\diamond + \psi(a^\diamond),$$

which amounts to $m < d - c$. As before, it then follows that

$$F_{\{\mathcal{C}^r \cup \mathcal{A}^r\}}(R, M) > F_{\{\mathcal{C}^r \cup \mathcal{A}^r\}}(M, M).$$

Therefore, Reg strictly dominates Mm . \square

Chapter 8

A Theory of Rationality

*Fra gli errori ci sono quelli che puzzano di fogna,
e quelli che odorano di bucato.*¹ (C. Pavese)

All theories of rational choice proposed so far in the literature have proven to have shortcomings. It is apparently in the nature of human behavior the impossibility of being fully described by a system of axioms and captured by some formal rules. As Amos Tversky used to say ([Gilboa, 2010]), “Give me an axiom and I’ll design the experiment that refutes it.” Our models and proposals in this and the previous chapters are no exception.

In this final part, we would nonetheless try to elaborate on the results of the preceding chapters, and to discuss a possible theory of rational choice.

Let us start off by reconsidering an example that was introduced in Chapter 3.

8.1 The Traveler’s Dilemma Case

An interesting examination of the Traveler’s dilemma has appeared in a paper by Ariel Rubinstein, entitled *Dilemmas of an Economic Theorist* ([Rubinstein, 2006]). In that paper, Rubinstein shows experimental results on the Traveler’s dilemma which he collected in various lectures at different universities during the years 2002-2003. The results of the experiments are reported in the following table.²

180	181-294	295	296-298	299	300
13%	15%	5%	3%	9%	56%

Rubinstein’s interpretation of the results is that

¹*Among errors, there are those that stink of sewer, and those that smell of fresh laundry.*

²In Rubinstein’s version of the game, players were allowed to claim an amount between \$180 and \$300.

The players who chose 180 are probably aware of the game theoretical prediction. On average, they would do badly playing against a player chosen randomly from the respondents. These players can claim to be the “victims” of game theory. The subjects whose answers were in the range 295–299 clearly exhibit strategic reasoning. The answer 300 seems to be an instinctive response in this context and the responses in the range 181–294 appear to be the result of random choice.

This interpretation is further corroborated by data about the subjects’ response time, showing that responses between 295 and 299 were also the slowest. Furthermore, he stresses:

Note that this regularity was found without any preconceived model and I am not aware of any existing game theoretical model that can, in fact, explain it.

Although Rubinstein wants to suggest that economists do not need to have a model in order to find regularities in the data, when looking at the results of the experiments a theorist would necessarily ask himself what models of rational choice could explain these regularities, and what kind of reasoning could be the basis of the behavior in the Traveler’s dilemma.

An answer to this question has been provided a few years later by Joe Halpern and Rafael Pass in a paper called *Iterated Regret Minimization: a New Solution Concept* ([Halpern and Pass, 2012]). They show that if players perform one round of elimination of actions that are dominated in terms of regret, then they are left with the set of actions $\{296, \dots, 300\}$. If a second round of elimination is performed on the remaining actions, only action 297 survives. This analysis in terms of regret is noticeably close to Rubinstein’s intuitions. The strategic reasoning acknowledged by Rubinstein may be based on, or at least explained by, regret minimization principles.

8.2 A Modest Proposal

Notice that in one of the quotes above, Rubinstein himself recognizes that the players acting in accordance with the standard game-theoretic analysis (called the “victims of game theory”) would do badly when randomly paired with a member of the population. This observation is perfectly in line with the evolutionary findings of Chapter 5 and Chapter 7.

However, the example of the Traveler’s dilemma is not meant to suggest that regret minimization is the choice principle that players *use*, just as the results of Chapter 5 and Chapter 7 are not intended to suggest that regret minimization is the criterion that players *should* use. The evolutionary advantage of regret minimization was established for a general class of environments, but not for all

possible environments that can possibly exist in nature. Other principles might be more beneficial for different environments.

These considerations lead us to some tentative guiding lights for a possible theory of rational choice.

1. Agents do not resort to a single and fixed principle in all decision situations.
2. The qualities of a choice cannot be evaluated independently of the environment in which it takes place.

Although these statements may sound trivial, their acceptance is not an uncontroversial issue.

Point 2, for example, suggests an ecological evaluation of choices, and it would be objected by the sustainers of a *subjectivist*, or *solipsistic*, approach to rationality (e.g., [Gilboa, 2014]). Such an approach would maintain that a choice is rational if DM does not feel embarrassed by it. Precisely,

According to this view, a mode of behavior is irrational for a decision maker, if, when the latter is exposed to the analysis of her choices, she feels uneasy or embarrassed by them. [...] Finally, *rationality* is defined by the negation of irrationality. Thus, a decision is rational for a decision maker if analysis of the decision, which could have been carried out by the decision maker at the time of decision, does not make one regret it.

The subjectivist approach delivers a weak notion of rationality, where the only requirement for a choice to be rational is to be accepted as such by the decision maker after further analyses of the decision problem. If DM will not feel embarrassed and will not reject the choice she made, then the choice is rational.

This is certainly a consistent definition of rationality, but it also displays some drawbacks, which the supporters of the subjectivist view are fully aware of. In Gilboa's words [Gilboa, 2014],

First, as opposed to behavioral axioms such as Savage's, this definition makes use of non-behavioral data. It does not suffice to know how a person behaves in order to determine whether they are rational. Rather, we need to find out whether they are embarrassed by their behavior. It is not clear how one can measure this embarrassment or unease, whether the expression of such emotions can be manipulated, and so forth.

Second, according to this definition the choice of axioms that define rationality cannot be made by decision theorists proving theorems on a whiteboard. Rather, the selection of axioms that constitute rationality becomes a subjective and empirical question: some people

may be embarrassed by violating an axiom such as transitivity, while others may not. [...]

Third, this definition also makes rationality non-monotonic in intelligence: suppose that two people make identical decisions, and that they violate a certain axiom. They are then exposed to the analysis of their decisions. One is bright enough to understand the logic of the axiom while the other isn't. The bright one, by feeling embarrassed, would admit that she had been irrational. By contrast, if the less intelligent person fails to see the logic of the axiom, he won't be embarrassed by violating it, and will be considered rational.

However, there must be reasons to embrace such a subjectivist approach, despite of the weaknesses that have been identified so far. According to Gilboa, the reason lies in the current state of the art in decision theory. As we have partly seen in Chapter 2, there is

conflict between the remarkable intellectual edifice of the rational choice paradigm and the vast body of experimental findings about violations of choice theoretic principles. [Gilboa, 2014]

We are then at a crossroad between two possible options to solve this conflict. Either we try to modify our theories by incorporating empirical findings in order to develop new theories that are descriptively more accurate, or we expose people to rational choice theories more extensively, hoping that this will make them behave less "irrationally". The advantage of the subjectivist position would then be that it can guide us in choosing a good direction at the crossroad we are now.

It is claimed that the definition of rationality suggested here offers the appropriate test for guiding us in this choice. If it is the case that most people who violate the theory are embarrassed by realizing how they behave, that is, if it is irrational for them to violate the theory, then it makes sense to teach the theory to them, and to hope that they will make better decisions in the future, according to their own judgment. If, by contrast, most people seem to be unperturbed by their violation of the theory, that is, it is rational for them to violate it, then there's little hope for the theory to be successful as a normative one, and we should accept people's behavior as a fact that's here to stay, and that should be incorporated into our descriptive theories to improve their accuracy. [Gilboa, 2014]

The problem we see with this motivation is that it is not clear, at least to us, what the rational choice theory that one should teach is supposed to be. What is the theory whose violations are supposed to embarrass people, and whose violations are hence supposed to be irrational? Classically, it would correspond to Savage's

axioms, but it is unlikely that Gilboa would refer to Savage's theory, since in the paper (as well as in other articles, see Chapter 2) Gilboa is ultimately arguing against the necessity of a Bayesian foundation of rational choice.

We are thus left with the previous question: what is the theory of rational choice that should be preached? As we have seen in Chapter 2, there are many competitors on the market, and there are consequently different interpretations of Gilboa's words. One possibility would be to expose the subjects to all the possible alternative theories. Once this has been done, we could either expect that all the subjects would decide to behave according to the same theory, or that each subject would pick his or her own favorite theory and always behave in accordance with it in all future circumstances, or neither of the two, and most people will continue to behave as before, in violation of each of the theories that have been taught. A further possibility is that people would decide to act according to different theories in different contexts. Since the first two possibilities seem highly unrealistic to us, we opt for the last one. Specifically, we would try to hold that in general people may follow different choice principles, depending on the decision situation at hand. If this is the case, then decision makers would irremediably exhibit violations of *the* theory, whatever the theory is. Decision makers do not stick to a single and fixed decision criterion for all their life, but are able to switch from one to another. Starting from a different approach, Gilboa, Postlewaite and Schmeidler have also reached similar conclusion in a different paper ([Gilboa et al., 2009]):

We reject the view that rationality is a clear-cut, binary notion that can be defined by a simple set of rules or axioms. There are various ingredients to rational choice. Some are of internal coherence, as captured by Savage's axioms. Others have to do with external coherence with data and scientific reasoning. The question we should ask is not whether a particular decision is rational or not, but rather, whether a particular decision is more rational than another. And we should be prepared to have conflicts between the different demands of rationality. When such conflicts arise, compromises are called for. Sometimes we may relax our demands of internal consistency; at other times we may lower our standards of justifications for choices. But the quest for a single set of rules that will universally define the rational choice is misguided.

For concreteness, going back to the framework of the previous chapters, we might expect agents to "use" maxmin expected utility in some cases and regret minimization in others. Obviously, whenever they will use maxmin expected utility they will possibly violate the axioms of regret minimization, and vice versa, when they will use regret minimization they will violate maxmin axioms.

From these premises, the obvious question to answer would be: what makes DM switch from one principle to another? Admittedly, what we are suggesting is far from being a fully-fledged theory, but we can at least propose a tentative

answer according to the guiding lights above. The switch might then be triggered by the *context* of the specific decision situation faced by the agent, in a way that has possibly been selected by evolution.

Context dependency is the key factor in our proposal, in two distinct and different uses of the term. Context dependency, as generally understood in decision theory, is the dependency on the menu of alternative options. This is the kind of context dependency that may generate intransitive choice patterns, and that calls for representation theorems in terms of general choice correspondences, as we have seen for the case of regret in Chapter 2. The second type of context dependency refers instead to the dependency on the environment in which choices take place, that is, the class of possible decision problems and interactions, and the composition and nature of the agents in the population. The analysis of the environment can improve the agents' decision making, and can show that certain context-dependent principles can be advantageous, even for rather general circumstances.

To be clear, we do not want to claim that agents are able to deterministically associate decision criteria with specific decision situations in an optimal way. Given the enormous number and unpredictability of possible choice situations that an individual might face during her life, it is likely that the switch would not be triggered by each single problem, but rather by some features that are shared by subsets of possible choice situations.³

Either way, given a certain environment, different associations of (features of) possible decision problems with choice principles will prove to be more or less beneficial from an evolutionary point of view, and successful associations might be selected accordingly. These considerations could then represent the conceptual framework for a theory of rational choice in line with the two points listed above, and with the general spirit of this work.

To conclude this section, notice that different modes of interaction between principles would be compatible with such a theory. Two principles might be competing, but they could also be combined, by weighting or nesting them, to obtain a final decision. The sense of this last possibility should become clearer in the next sections.

8.3 Relief Maximization

A popular book by Dan Ariely [Ariely, 2008], entitled *Predictably Irrational*, begins with the case of an advertisement for a yearly subscription to the *Economist*. The customer was offered the following three alternatives:

1. online subscription, \$59;

³This perspective can also be related to considerations about cognitive limitations (see [Zollman, 2008], [Bednar and Page, 2007], [Mengel, 2012], [Rayo and Becker, 2007]).

2. print subscription, \$125;
3. print-and-online subscription, \$125.

The last two options are both priced at \$125, there is no typo. The second option is obviously dominated by the third. When Ariely wanted to test the effect of including the dominated alternative (called “decoy option”), he asked one hundred MIT students to choose one of the options in the menu and found that the students decided as follows:

1. 16 students chose the online subscription;
2. 0 students chose the print subscription;
3. 84 students chose the print-and-online subscription.

Then Ariely removed the decoy option from the menu, and asked the students again. He discovered that in this case 68 students opted for the online subscription, while only 32 chose the print-and-online subscription. The mere existence of a dominated option, that nobody opted for, caused a substantial difference in the choice between the other two alternatives. The first chapter of his book is full of similar examples, where the decoy effect is evidently playing a considerable role in the individuals’ decisions. The general observation made by Ariely is that

[...] humans rarely choose things in absolute terms. We don’t have an internal value meter that tells us how much things are worth. Rather, we focus on the relative advantage of one thing over another, and estimate value accordingly.

Relativity is (relatively) easy to understand. But there’s one aspect of relativity that consistently trips us up. It’s this: we not only tend to compare things with one another but also tend to focus on comparing things that are easily comparable—and avoid comparing things that cannot be compared easily.

These ideas show in the subscription case in that the decision is structured in a way that makes the comparison between the second and third option straightforward, while it is not obvious how the first option compares with these two. The *Economist’s* marketing wizards aim at exploiting this asymmetry in order to sell the subscription they want to sell. And it works: when the second option is present, the third option apparently looks much better.

The case presented by Ariely is an example of decision under certainty, while all this study was focused on decisions under uncertainty. Decoy effects, also called asymmetric dominance effects, have been recently investigated in the context of decisions under uncertainty, and in game theory in particular. In [Colman et al., 2007], the authors conducted an experimental analysis on a set of two-player games with

three possible actions, one of which (the decoy action E) is dominated by one of the remaining two (action C), and with no other dominance relation between the other actions. Their experiments were performed both on symmetric and asymmetric games, with similar results. Given the single-population models we have worked with so far, we are mainly interested in the symmetric case. The symmetric games used in the experiments are the following.⁴

I	<table border="1"><tr><td></td><td>C</td><td>D</td><td>E</td></tr><tr><td>C</td><td>40</td><td>20</td><td>60</td></tr><tr><td>D</td><td>60</td><td>40</td><td>20</td></tr><tr><td>E</td><td>20</td><td>0</td><td>40</td></tr></table>		C	D	E	C	40	20	60	D	60	40	20	E	20	0	40
	C	D	E														
C	40	20	60														
D	60	40	20														
E	20	0	40														

II	<table border="1"><tr><td></td><td>C</td><td>D</td><td>E</td></tr><tr><td>C</td><td>60</td><td>20</td><td>40</td></tr><tr><td>D</td><td>0</td><td>80</td><td>0</td></tr><tr><td>E</td><td>20</td><td>0</td><td>20</td></tr></table>		C	D	E	C	60	20	40	D	0	80	0	E	20	0	20
	C	D	E														
C	60	20	40														
D	0	80	0														
E	20	0	20														

III	<table border="1"><tr><td></td><td>C</td><td>D</td><td>E</td></tr><tr><td>C</td><td>80</td><td>20</td><td>40</td></tr><tr><td>D</td><td>40</td><td>80</td><td>0</td></tr><tr><td>E</td><td>60</td><td>0</td><td>20</td></tr></table>		C	D	E	C	80	20	40	D	40	80	0	E	60	0	20
	C	D	E														
C	80	20	40														
D	40	80	0														
E	60	0	20														

IV	<table border="1"><tr><td></td><td>C</td><td>D</td><td>E</td></tr><tr><td>C</td><td>60</td><td>20</td><td>40</td></tr><tr><td>D</td><td>0</td><td>80</td><td>0</td></tr><tr><td>E</td><td>40</td><td>0</td><td>20</td></tr></table>		C	D	E	C	60	20	40	D	0	80	0	E	40	0	20
	C	D	E														
C	60	20	40														
D	0	80	0														
E	40	0	20														

V	<table border="1"><tr><td></td><td>C</td><td>D</td><td>E</td></tr><tr><td>C</td><td>80</td><td>20</td><td>20</td></tr><tr><td>D</td><td>40</td><td>80</td><td>0</td></tr><tr><td>E</td><td>60</td><td>0</td><td>0</td></tr></table>		C	D	E	C	80	20	20	D	40	80	0	E	60	0	0
	C	D	E														
C	80	20	20														
D	40	80	0														
E	60	0	0														

Players first had to play the 3×3 games, and later, as control condition, they had to play the same games except that the dominated action E had been deleted. In short, the findings were that asymmetric dominance effects were significantly exhibited in most of the games. Games II and IV represented natural exceptions, in the authors' opinion, because of the payoff dominance of the (D, D) equilibrium:

[...] the payoff dominance of the (D, D) outcome may have counteracted the influence of the strategic dominance of C over E to some extent in the 3×3 versions of these games by providing players with a persuasive reason for choosing D . [...] The asymmetric dominance effect seems most likely to emerge when the control version of a game lacks a focal point on which players might expect to coordinate, presumably because asymmetric dominance provides a reason for choice, but it can be overwhelmed by a strong focal point such as a payoff-dominant equilibrium that offers an alternative reason for choice.

The last quote underlines that asymmetric dominance may provide a “reason for choice”, which sounds something very similar to a choice principle. Interestingly enough, we can formally express this reason for choice as a decision criterion like those listed in Chapter 5:

- **Simple Relief Maximization:**

$$\hat{a}(u, B) = \operatorname{argmax}_{a \in A} \min_{b \in B} \{u(\pi(a, b)) - \min_{a' \in A} u(\pi(a', b))\}.$$

⁴The symmetric games presented in [Colman et al., 2007] are actually six, but one of them is uninteresting for our purposes, because it does not distinguish between any of the choice principles considered here. As usual, since games are symmetric, it suffices to specify row player's payoffs.

In parallel to the probabilistic version of maxmin and regret minimization, it is also straightforward to define the probabilistic version of relief maximization:

- **Expected Relief Maximization:**

$$\hat{a}(u, \Gamma) = \operatorname{argmax}_{a \in A} \min_{P \in \Gamma} \{ \mathbb{E}_P [u(\pi(a, b))] - \min_{a' \in A} \mathbb{E}_P [u(\pi(a', b))] \}.$$

To the best of our knowledge, the concept of relief has been formally introduced in decision theory by Richard Jeffrey in [Jeffrey, 1990], and has gone unused since then. On the other hand, the link between relief maximization and decoy effects is apparent: the choice that would be determined by asymmetric dominance effects is precisely the choice that maximizes simple relief in all of the above examples.

Is acting according to asymmetric dominance irrational? Well, first of all, it can be easily checked that both simple and expected relief maximization violate the independence of irrelevant alternatives, and that consequently both violate Savage's axiom P1 (see Chapter 2).

Second, apparently decoy effects are not held in the highest regard by behavioral economists either. In this respect, Ariely writes

This is not only irrational but predictably irrational as well.

According to the subjectivist view, however, it might also be fully rational to act in accordance to relief maximization. Suppose a person exhibits the decoy effect in the subscription example above, i.e., she would choose the print-and-online option when the dominated option is also available, and the online subscription otherwise. Shall this person feel embarrassed by her choices? It is indeed possible that she might feel embarrassed by violating the independence of irrelevant alternatives, but it is also possible that she would defend her actions in light of the reason for choice provided by relief-maximizing considerations. The same would hold for a game-theoretic context too. In the presence of the dominated option, DM has a reason for defending her choice, and a possible choice principle that she could appeal to. From a subjectivist perspective, asymmetric dominance would be irrational for a decision maker that acts, or wants to act, in accordance to Savage's axioms (or, equivalently, maximization of expected utility), but it would be perfectly rational for an agent that wants to maximize relief.⁵ From this point of view, asymmetric dominance effects may or may not be rational, depending on the choice principle motivating DM's choices.

According to point 2 of our tentative guiding lights, however, the qualities and the rationality of a choice cannot be evaluated independently of the environment in which it takes place. This approach would then recommend an ecological and

⁵The same would hold for regret minimization. As a personal consideration, when I look, for example, at Stoye's axiomatization of regret minimization as opposed to Savage's axioms, I don't feel disgusted by Stoye's axioms. On both sides, there are axioms that seem to me very reasonable, and others less so.

evolutionary analysis of relief maximization as a choice principle. What would evolutionary selection have to say about relief maximization?

8.4 From an Evolutionary Point of View

Relief maximization as choice principle is rather weak. Its weakness lies in the fact that relief has no bite unless there is in the decision problem an action that is a never-worst-reply. In problems without such actions, simple relief maximization would be indifferent between all alternatives available. For a decision criterion, this is intuitively a clear shortcoming, that should somehow manifest also in evolutionary contexts and contests.

Consider, for instance, the game of life of Chapter 5 where all games are 2×2 games. In such an environment, relief maximizers will perform very poorly, both against maximinimizers and against regret minimizers. Indeed, relief maximizers would be indifferent (and, consequently, would choose at random) in all games apart from those where one of the two actions strictly dominates the other, which is also the behavior prescribed by both maxmin and regret minimization in all games with a strictly dominant action.

In general, relief maximizers would lose the evolutionary competition against both maximinimizers and regret minimizers in many environments that are not explicitly tailored to relief maximization.

To have an idea of this, consider an environment consisting of the five games of the previous section, and suppose for the sake of the example that each of them has the same occurrence probability. Suppose, moreover, that the only two types in the population are the maximinimizers and the relief maximizers. The expected fitness in the corresponding game of life is given by the next table (in the row player's perspective).

	$Mm, \tilde{\Gamma}$	$Rel, \tilde{\Gamma}$
$Mm, \tilde{\Gamma}$	64	66
$Rel, \tilde{\Gamma}$	62	64

There, relief maximization (and, hence, asymmetric dominance effects) turns out to be strictly dominated by maxmin as a choice principle. The only evolutionarily stable state will then be a monomorphic population of maxmin players.

One might now wonder why asymmetric dominance effects are still consistently observed in actual behavior, given that agents who fall prey to decoys should go extinct in many circumstances. There are two possible explanations. The first is that relief-based reasoning survived as a side effect of the evolutionary success of other context-dependent principles, such as regret minimization. In the case of regret minimization, context dependency is the key of its evolutionary advantage. It is possible that human beings generally retained a propensity for choosing in a context-dependent manner (by comparing options with each other),

that is evolutionarily beneficial when implemented in terms of regret, but much less beneficial when implemented in terms of relief. This direction would be in line with the previous observation by Ariely that “humans rarely choose things in absolute terms.” The human attitude of judging things relative to the context might then be much more rational than it is normally believed in the literature. Not all context-dependent principles are equally good, though.

A second option is to ponder on the possibility that different choice principles could also be mixed. Point 1 of our guiding lights excludes the idea that agents use the same decision criterion in all possible circumstances. As outlined before, different principles may then be triggered by different (features of) choice situations. But agents might also combine, in a lexicographic or in a weighted way, different principles in the same decision problem.

8.5 Nested and Weighted Principles

It might be the case in some decision situations that two or more actions are equally optimal according to a certain decision criterion. How will DM break symmetry and choose between them? In Chapter 5, we assumed that DM will simply pick one at random, but an alternative would be to reconsider the equally optimal actions in the light of a different principle, i.e., to use a second principle to break the tie.

Let us consider the previous five games again.⁶ But notice that, since in all those games the Nash equilibria are in the diagonal, they somehow provide an incentive for coordination between the players. To have a more diverse environment, let us also consider their “anti-coordination” versions, where the first two rows are swapped around.

I'	<i>D</i>	<i>C</i>	<i>E</i>
<i>D</i>	60	40	20
<i>C</i>	40	20	60
<i>E</i>	20	0	40

II'	<i>D</i>	<i>C</i>	<i>E</i>
<i>D</i>	0	80	0
<i>C</i>	60	20	40
<i>E</i>	20	0	20

III'	<i>D</i>	<i>C</i>	<i>E</i>
<i>D</i>	40	80	0
<i>C</i>	80	20	40
<i>E</i>	60	0	20

IV'	<i>D</i>	<i>C</i>	<i>E</i>
<i>D</i>	0	80	0
<i>C</i>	60	20	40
<i>E</i>	40	0	20

V'	<i>D</i>	<i>C</i>	<i>E</i>
<i>D</i>	40	80	0
<i>C</i>	80	20	20
<i>E</i>	60	0	0

Two remarks are in order here. Firstly, notice that *E* is still dominated by *C*, and no other dominance relations exist in the games, as before. So *C* is still the relief maximizing action in all games. Secondly, with respect to the description of their experiment, [Colman et al., 2007] say:

⁶Remember that we didn't choose those games. The selection comes from [Colman et al., 2007].

The dominant strategies are shown in row and column C and the dominated strategies in row and column E , although in the experiment, to control for positioning and labeling effects, strategies were rotated systematically so that they appeared equally frequently in all three rows and columns.

The authors seem to assume that the rotation of actions in the matrix does not affect the players' choices, but this is not obvious at all to us. Consider games II and II' for example. In game II, (C, C) and (D, D) are Nash equilibria of the game, while in game II' the Nash equilibria are (C, D) and (D, C) . The coordination problem of game II looks easier to solve than the anti-coordination problem of game II'. In particular, only in game II the authors can appeal to the payoff-dominance justification for equilibrium (D, D) , but not in game II'. However, they do not seem to distinguish between game II and its rotated version II' in their analysis.

Either way, let us get back to our evolutionary analysis. Suppose a population of maximinimizers and regret minimizers is surrounded by an environment composed of the previous ten games, and each game has the same probability of occurring. The expected fitness of the two types in the corresponding game of life is given in the following table.

	$Mm, \tilde{\Gamma}$	$Reg, \tilde{\Gamma}$
$Mm, \tilde{\Gamma}$	44	44
$Reg, \tilde{\Gamma}$	46	46

This is another instance of what we have seen in Chapter 5: maxmin is strictly dominated by regret minimization. Let us now introduce two more sophisticated variants of maximinimizers and regret minimizers, that first evaluate choices according to one of the two principles, and then proceed to a second round of evaluation of the surviving actions in terms of relief. So, type Mm^+ first discards all actions that do not maximize the minimum, and then, among the surviving actions, discards all those that do not maximize relief. The resulting expected fitness of the four types is specified by the next table.

	$Mm, \tilde{\Gamma}$	$Mm^+, \tilde{\Gamma}$	$Reg, \tilde{\Gamma}$	$Reg^+, \tilde{\Gamma}$
$Mm, \tilde{\Gamma}$	44	44	44	44
$Mm^+, \tilde{\Gamma}$	42	42	42	42
$Reg, \tilde{\Gamma}$	46	46	46	46
$Reg^+, \tilde{\Gamma}$	46	46	46	46

Reasoning in terms of relief, in combination with other principles, is not necessarily detrimental to the evolutionary fitness. When paired with maxmin, relief minimization performs rather poorly, but when paired with the other context-dependent principle, regret minimization, relief reasoning can survive the evolutionary competition. In the example considered, the evolutionarily stable states are all population states with types Reg and Reg^+ only.

8.6 The Power of Context Dependency

The fact that humans do not usually choose in absolute terms, but rather evaluate options in comparison with each other, is not a new discovery, and context-dependent choices have been extensively studied in behavioral economics (see, among others, [Kahneman and Tversky, 1979], [Tversky and Kahneman, 1974], [Tversky and Kahneman, 1981], [Tversky and Simonson, 1993], [Shafir et al., 1993]). However, context-dependent effects have never had a highly regarded reputation in decision theory. They are normally viewed as a consequence of human imperfect decision making, and are deemed as irrational phenomena, maybe predictably irrational. In this respect, [Tversky and Kahneman, 1981] for example write:

The definition of rationality has been much debated, but there is general agreement that rational choices should satisfy some elementary requirements of consistency and coherence.

[...] Because of imperfections of human perception and decision, however, changes of perspective often reverse the relative apparent size of objects and the relative desirability of options.

A famous case study that Tversky and Kahneman present in their paper is about two alternative programs to combat a disease which is expected to kill 600 people. The two alternatives were given as follows:

- if Program A is adopted, then 200 people will be saved;
- if Program B is adopted instead, then there is a probability of $\frac{1}{3}$ that 600 people will be saved, and a probability of $\frac{2}{3}$ that nobody will be saved.

When they asked a first group of participants in the experiment what they would have opted for, 72% of the participants chose Program A, and 28% chose Program B. Next, they presented a second group of participants the same problem, except that it was phrased as follows:

- if Program C is adopted, then 400 people will die;
- if Program D is adopted instead, then there is a probability of $\frac{1}{3}$ that nobody will die, and a probability of $\frac{2}{3}$ that 600 people will die.

Most of the participants from the second group opted for Program D (78%), while only 22% chose Program C.

When we require coherence and consistency as necessary attributes of rationality, the results of the experiment cannot be justified as rational. As they notice in the paper, the framing effect of phrasing the problem in terms of losses (deaths) rather than gains (survivals) should not affect the people's choice according to standard models of rationality. As a side comment, to check which of the two aspects prevails, it would be interesting to test people's reaction to a third version of the problem:

- if Program E is adopted, then 400 people will die and 200 will be saved;
- if Program F is adopted instead, then there is a probability of $\frac{1}{3}$ that nobody will die and everyone will be saved, and a probability of $\frac{2}{3}$ that 600 people will die and nobody will be saved.

Unfortunately, this experiment is not reported in [Tversky and Kahneman, 1981]. You can still ask yourself what you would choose in this case.

In general, Tversky and Kahneman describe these phenomena as “shifts of preference”. Differently from the evolution of preference literature that we have seen in earlier chapters, by preference they mean DM’s observed action choices, and not the subjective utilities over outcomes. But what the two different formulations are suggesting more specifically is, perhaps, a shift of decision criterion. Again, rather than a general theory, the reader should consider this just as a suggestion for a possible interpretation of the experiment to be further tested in the future.

The first formulation of the choice situation would then make DM focus on the perspective of saving at least 200 persons. The decision problem could thus be represented as pictured in the following table.

	$\frac{1}{3}$	$\frac{2}{3}$
A	200	200
B	600	0

Instead, the second formulation of the problem could trigger a more regret-based perspective on the problem, and suggest a reading of the situation possibly close to the next table.

	$\frac{1}{3}$	$\frac{2}{3}$
C	-400	0
D	0	-200

It is possible that DM, when faced with questions of life and death like this, could have a propensity for using a security strategy, i.e., some security reasoning based on worst-case considerations, such as simple maxmin, simple regret minimization, or simple relief maximization. But then the best worst-case scenario (and the corresponding action) would be different, depending on the specific security principle used. According to this possibility, the first formulation would trigger the use of maxmin principles, while the second formulation would favor regret-minimizing considerations.

One of the main points stemming from the results in previous chapters is that context dependency should not necessarily be viewed as irrational. Regret-minimizing preferences are context-dependent, but nevertheless are able to outperform context-independent choice principles, like maxmin, in considerably general circumstances. Among context-dependent criteria, though, adopting the appropriate perspective when comparing available options is crucial. Some perspectives are more beneficial than others, as highlighted by the examples on relief.

Nonetheless it would be better, at least from an evolutionary point of view, to distinguish between different context dependencies rather than rejecting them as a whole in the name of some essential axiom of rationality.

Clearly, we do not adopt the approach of judging what is rational (exclusively) based on the acceptability, or acceptance, of a system of axioms. This is mainly because of two sorts of considerations. First of all, we do not want to constrain the agent's decision making to any fixed decision criterion. We believe it is common and perfectly admissible to shift from a decision criterion to another given different circumstances, so that we cannot link the rationality of an agent to any fixed axiomatic system.

The second consideration is related to evolutionary arguments. Suppose there is a large population of maximinimizers and regret minimizers facing a set of possible interactions (and assume for simplicity that the games that are played in the population have monetary payoffs, and that fitness corresponds to money). Suppose there is an external observer, an analyst named Charles, who has two good friends in the population, Ann and Bob. As an analyst, Charles decides to analyze the situation. Given the relevant information available, i.e., given the games that might possibly occur and given what has been observed about the composition of the population so far, he realizes that choosing according to regret minimization could actually be beneficial. Charles shows his two friends the models that led him to this conclusion (models like the games of life introduced before for example, or something else), and Ann and Bob agree that the assumptions about the environment and the population composition in Charles' models are reasonable and in line with their knowledge about the situation. Finally, after listening to his arguments, they are persuaded that Charles has good reasons for arguing in favor of regret minimization. But then their common friend David comes along, and strongly discourages Ann and Bob to act in accordance with regret minimization, because it is not transitive, it can induce preference cycles, at the risk of being money-pumped, etc. David manages to convince Bob to avoid regret minimization and to use, say, maxmin instead, while Ann decides to stick to regret minimization. David's analysis is surely correct, regret-based principles can create cycles, and can be money-pumped, while maxmin generates a preference ranking of options that is independent of the available alternatives. But, for the sake of the argument, let us suppose that Charles' analysis was also correct: given the environment, regret minimization turns out to be more fruitful than maxmin. Consequently, what will happen is that Bob will act according to maxmin and avoid possibly intransitive choices, while Ann will choose according to regret minimization and will be richer at the end of the day. Now, if Bob primarily cares about money, *shouldn't* he shift to regret minimization after realizing that Charles was right in his analysis?

These examples lead us to considering which arguments are normatively more powerful. If both Ann and Bob primarily cared about their personal wealth (or evolutionary fitness), would the axiomatic argument be more compelling than

the evolutionary argument from a normative point of view? Ann is becoming richer, while Bob is firmly avoiding any possible violation of transitivity: is it then more likely that Ann will be persuaded that she should start using maxmin, or that Bob will be persuaded that he should start using regret minimization? The second, in our opinion. Or at least, if we didn't observe Bob switching to regret minimization, we might have to assume that what he primarily cares about is transitivity, and not final wealth or fitness.

Suppose that there is a change in the environment, and according to Charles' analysis maxmin will now bear greater wealth. Between maxmin and regret minimization, Charles then suggests his friends Ann and Bob to play according to maxmin. Of course Charles can be wrong, but he was right in his previous suggestion, and let us suppose that Charles' analysis is again correct. Shouldn't Ann and Bob switch to the maxmin principle then? Or shall they ponder over different axiomatic systems and opt for the choice principle with nicer axioms? Assuming that the environment changed and that Charles is right, we maintain that they should switch. In other words, if agents care about the final amount of money, we would suggest to act in accordance with regret-minimizing reasoning when playing the Traveler's dilemma, but to adapt to transitive criteria, such as maxmin, in the presence of possible money pumps. Would it be irrational to change from a criterion to another depending on the structure of the decision problem and on the environment?

Notice, however, that the appeal to evolutionary fitness is not even necessary for the previous argument.⁷ Consider for example the results of Chapter 5, and suppose that the function Φ is no longer the fitness function, but just a subjective utility over outcomes. In such a situation we would have a population of agents with the same subjective utility and different choice principles. Even if a player were only interested in the amount of personal happiness Φ achieved at the end of the day, then Charles would advise to adopt regret minimization instead of maxmin as decision criterion in the context of the game of life defined there. Indeed, for the circumstances described in Chapter 5, the expected subjective happiness (independent of the evolutionary fitness) of a regret minimizer would be superior than that of a maximinimizer.

As a final remark, a parallel and complementary research direction is to investigate the evolution of theory of mind. Theory of mind (ToM) is the ability of attributing mental states, such as beliefs and desires, to other agents, and it is an essential part of strategic reasoning. Agents with ToM level-0 act as if they have no beliefs or desires, level-1 agents instead form a belief over the co-player's actions and choose a best reply to that belief, and hence they will avoid strictly dominated actions. Level-2 agents assume that the co-player is of level-1 and does not play dominated actions, and therefore they will only play actions that survive two iterations of elimination of strictly dominated actions. Level-3 agents

⁷I am indebted to Tommaso Orlando for this observation.

think that the co-player is a level-2 agent, and so on (see [Camerer et al., 2004], [Ho et al., 1998], [Stahl and Wilson, 1994], [Stahl and Wilson, 1995]).

Apparently, humans did not develop the ability of iterating more than 2-3 times the elimination of dominated actions (e.g., [Beard and Beil, 1994] and [Ho et al., 1998]). This might have two explanations: either they cannot reason higher than 2-3 levels of theory of mind, or they don't believe that the co-player can reason more than 1-2 levels. Both explanations may be supported by ecological arguments. [Stahl, 1993] for example shows that the survival of super-intelligent players, with the ability of reasoning up to any level of theory of mind, can be questioned from an evolutionary perspective. If higher levels of reasoning are not costless, then evolution could select against super-smart agents. Other recent results suggest that reasoning more than one level higher than your co-player may be self-defeating in strategic interactions (see [Mohlin, 2012], [de Weerd et al., 2013]). The best position would be to outdo your co-player by exactly one level. Taken together, these results point again to the relevance of the environment: agents with the highest level of ToM would probably be outperformed by agents whose level is only one step above the majority of the population. Similar to the case of choice principles, the performance of a specific level of ToM is contextual and essentially depends on the environment.

8.7 Conclusion

We do not have conclusive evidence to support the intuitions advanced in this chapter, which are destined to simply remain a modest proposal for now. But we hope we have at least suggested some reasons to take into consideration, if only to reject it in the future, a more ecological and less axiomatic approach to the normative debate on the rationality of decision criteria.

With this perspective, we also share with Gilboa the wish to better integrate normative and descriptive accounts of rationality. Indeed, if previous considerations support a normative reading of this proposal, it is also reasonable to expect that successful principles (or successful associations between principles and contexts) had a higher chance to survive, while deleterious principles and unfortunate associations should have gone extinct. From a descriptive point of view, we could then presume that the observed behavior shall somehow reflect the results of such a natural selection. By studying the contexts and environments in which agents developed their decision making, we may hope to also improve the descriptive accuracy of our theories of rational choice.

Despite the scarce interest and low regard demonstrated towards context dependency, in the sense of dependency of action choices on both the menu of available options and on the environment, the success of an agent's decisions seems to be essentially affected by these two dimensions of contextuality. Choice principles that evaluate options in comparative rather than absolute terms can be favored

by evolutionary selection in many general circumstances. The appeal to different principles on the basis of the analysis of the specific environment and decision situation may also be beneficial to the agent's decision making. Experimental results at odds with standard theories of rational choice might be better understood if interpreted as the behavior resulting from a more context-dependent decision making, that evolution could have positively selected, without any consideration about coherence, consistency, or acceptability of different systems of axioms.

When we started researching on the issue of rational choice, our aim was initially to find a better principle than Savage's maximization of subjective expected utility, which still constitutes the standard definition of rationality in economics. A predicament immediately encountered in pursuing our goals was that it is not even established what the proper ground is on which one should start building the very first notion of rationality. Should rationality be investigated from an axiomatic perspective, by proposing and debating over different systems of axioms? Should it be considered instead from an evolutionary point of view, in terms of the ecological rationality of different choices? Does it appertain to psychology, in the sense of describing the rationale behind the agent's decisions? We straight-away found ourselves in an impasse: different bases have been proposed for the foundation and the study of the concept of rationality without ever reaching a consensus on where one should even begin the investigation.

In Chapter 2, we presented and extensively discussed two of the possible approaches to rational choice. Since we think of the issue of rationality primarily as normative (as explained in Chapter 1), we left aside any consideration about the psychology of reasoning,¹ and we directed our attention to the dualism between the axiomatic approach and the ecological approach.

Savage's classic work obviously belongs to the axiomatic approach. His axioms, as we have seen, have been attacked mostly from two different perspectives, that can be called an axiomatic perspective, and a behavioral perspective. Axiomatic attacks came from those who tried to reject (mainly with normative arguments) one or more of Savage's axioms, and to propose alternative axioms for rationality. This is the case, among others, of [Gilboa and Schmeidler, 1989] and [Schmeidler, 1989]. According to this position, Savage's axioms and subjective expected utility maximization represent an unsatisfactory norm for rationality, and different systems of axioms and better decision criteria are needed. Behavioral attacks, on the other hand, are often descriptive, and criticize SEU theory because

¹But see also [Stenning and van Lambalgen, 2008] for a detailed study in this direction.

it is systematically at odds with decision makers' behavior. This is the case, for instance, of [Allais, 1953], [Ellsberg, 1961], and [Kahneman and Tversky, 1979].

Our approach is different. We didn't want to argue for a specific system of axioms over another. Neither did we want to object against some specific axioms, or choice principles, from a descriptive perspective, because they would be in contradiction with observed behavior.² We opted instead for a third, less well-trodden path: the study of the ecological rationality of different choice principles.

The idea of investigating the ecological rationality of choices traces back at least to the work by Herbert Simon, and has been further developed by Gigerenzer and colleagues (see Chapter 2). Here we proposed a general model, the game of life, to allow the ecological analysis and the evolutionary comparison of classic choice principles coming from the literature in decision theory. Maximization of subjective expected utility is just one among many, and, in a sense, a special case for situations where uncertainty is represented in a probabilistic way. One of the important results is then that if we don't buy the assumption that all uncertainty is and must be probabilistic, regret-based choice may prove to be evolutionarily successful and ecologically beneficial in various environments. Furthermore, as specified also in Chapter 5, even if we want to advocate that all uncertainty should be quantified in probabilistic terms,³ regret-based principles are never disfavored by evolutionary selection in all scenarios considered here.

In the end of all this research, however, we reached the conclusion that the goal we were aiming at when we started is actually misleading. Indeed, the notion of ecological rationality hinges on the specific environment that we take into account, and it is essentially relativistic in this respect. The same principle can be very good or very bad, evolutionarily speaking, depending on the surrounding environment, which consists of at least two aspects: the (structure of the) decision problems an agent has to face, and the composition of the population the agent is part of. From this perspective, the quality of a choice is fundamentally related to the environment, and not objectively given a priori. And so is the rationality of different choice principles considered here. The notion of a "better principle" is hence meaningless in itself, and it should rather be replaced by the more contextual notion of a "better principle, given a certain environment".

Finally, this work still leaves many open issues for further developments, such as the ones sketched in Chapter 8. One of those that we consider very interesting and urgent for this program is the investigation of possible shifts from a decision criterion to another. In our opinion, it is a fact that agents do not exclusively stick to the same principle in all decision problems throughout their life. Rather, they are able to switch from one criterion to another, and, we conjecture, the shift may be triggered by the context and the structure of the specific decision problem

²Here we just take as a matter of fact that SEU maximization is descriptively inaccurate as decision criterion.

³We already blatantly disagreed with this position all along.

in a way that reflects evolutionary success. This thesis only sketched a possible theory along these lines (see Chapter 8), and wishes to see more advancements in this direction.

Moreover, the model introduced here allows the comparative study of other features traditionally examined in decision theory, such as different attitudes towards risk and towards ambiguity, and different levels of theory of mind, as already anticipated in Chapter 8. The combination of all this factors and its evolutionary analysis might shed new light on how we (should) make choices in different contexts.

Bibliography

- [RLA, 1980] (1980). Right Livelihood award. <http://rightlivelihood.org/>.
- [GPP, 2016] (2016). Global Priorities Project. <http://globalprioritiesproject.org/2016/04/global-catastrophic-risks-2016/>.
- [Alchian, 1950] Alchian, A. (1950). Uncertainty, evolution and economic theory. *Journal of Political Economy*, 58:211–221.
- [Alger and Weibull, 2013] Alger, I. and Weibull, J. W. (2013). Homo Moralis: Preference Evolution Under Incomplete Information and Assortative Matching. *Econometrica*, 81:6:2269–2302.
- [Allais, 1953] Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: Critique des postulats et axiomes de l’ecole americaine. *Econometrica*, 21(4):503–546.
- [Anderson, 1991] Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14(3RA):471–517.
- [Anscombe and Aumann, 1963] Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1):199–205.
- [Ariely, 2008] Ariely, D. (2008). *Predictably Irrational*. HarperCollins.
- [Aumann, 1976] Aumann, R. (1976). Agreeing to Disagree. *Annals of Statistics*, 4:1236–1239.
- [Aumann, 1995] Aumann, R. (1995). “Interactive Epistemology”, Discussion Paper No. 67. Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem.

- [Aumann and Brandenburger, 1995] Aumann, R. and Brandenburger, A. (1995). Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63:1161–1180.
- [Baltag and Smets, 2008] Baltag, A. and Smets, S. (2008). A Qualitative Theory of Dynamic Interactive Belief Revision. *Texts in Logic and Games*, 3:9–58.
- [Baltag et al., 2009] Baltag, A., Smets, S., and Zvesper, J. (2009). Keep ‘Hoping’ for Rationality: A solution to the backward induction paradox. *Synthese*, 169:301–333.
- [Basu, 1994] Basu, K. (1994). The Traveler’s Dilemma: Paradoxes of Rationality in Game Theory. *The American Economic Review*, 84(2):391–395.
- [Bather, 1996] Bather, J. (1996). A conversation with herman chernoff. *Statist. Sci.*, 11(4):335–350.
- [Battigalli, 2016] Battigalli, P. (2016). Game Theory: Analysis of Strategic Thinking. Unpublished manuscript.
- [Battigalli et al., 2015] Battigalli, P., Cerreia-Vioglio, S., Maccheroni, F., and Marinacci, M. (2015). Self-confirming equilibrium and model uncertainty. *American Economic Review*, 105:2:646–677.
- [Battigalli and Siniscalchi, 2002] Battigalli, P. and Siniscalchi, M. (2002). Strong Belief and Forward Induction Reasoning. *Journal of Economic Theory*, 106:356–391.
- [Beach and Mitchell, 1978] Beach, L. R. and Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *The Academy of Management Review*, 3(3):439–449.
- [Beard and Beil, 1994] Beard, T. and Beil, R. (1994). Do People Rely on the Self-Interested Maximization of Others? An Experimental Test. *Management Science*, 40(2):252–262.
- [Becker, 1976] Becker, G. S. (1976). Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology. *Journal of Economic Literature*, 14:817–826.
- [Bednar and Page, 2007] Bednar, J. and Page, S. (2007). Can game(s) theory explain culture?: The emergence of cultural behavior within multiple games. *Rationality and Society*, 19(1):65–97.
- [Bell, 1982] Bell, D. E. (1982). Regret in Decision Making Under Uncertainty. *Operations Research*, 30:5:961–981.
- [Bergstrom and Godfrey-Smith, 1998] Bergstrom, C. T. and Godfrey-Smith, P. (1998). On the evolution of behavioral heterogeneity in individuals and populations. *Biology and Philosophy*, 13(2):205–231.

- [Bernoulli, 1738] Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*. translated as “Exposition of a new theory on the measurement of risk”, *Econometrica*, 22, 23–36, 1954.
- [Bester and Güth, 1998] Bester, H. and Güth, W. (1998). Is Altruism Evolutionarily Stable? *Journal of Economic Behavior and Organization*, 34:193–209.
- [Birnbaum, 2008] Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, 115:463–501.
- [Blackburn et al., 2001] Blackburn, P., de Rijke, M., and Venema, Y. (2001). *Modal Logic*. Cambridge University Press, Cambridge.
- [Bleichrodt and Wakker, 2015] Bleichrodt, H. and Wakker, P. P. (2015). Regret theory: A bold alternative to the alternatives. *The Economic Journal*, 125(583):493–532.
- [Brafman and Tennenholtz, 2000] Brafman, R. I. and Tennenholtz, M. (2000). An axiomatic treatment of three qualitative decision criteria. *J. ACM*, 47(3):452–482.
- [Brandenburger, 1992] Brandenburger, A. (1992). Knowledge and Equilibrium in Games. *Journal of Economic Perspectives*, 6:83–101.
- [Brandenburger, 2008] Brandenburger, A. (2008). Epistemic game theory: an overview. In Durlauf, S. and Blume, L., editors, *The New Palgrave Dictionary of Economics*. Palgrave MacMillan.
- [Brandenburger and Dekel, 1993] Brandenburger, A. and Dekel, E. (1993). Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189–198.
- [Camerer et al., 2004] Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898.
- [Charness and Rabin, 2002] Charness, G. and Rabin, M. (2002). Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics*, 117:3:817–869.
- [Chater and Oaksford, 2000] Chater, N. and Oaksford, M. (2000). The rational analysis of mind and behavior. *Synthese*, 122:93–131.
- [Colman et al., 2007] Colman, A. M., Pulford, B. D., and Bolger, F. (2007). Asymmetric dominance and phantom decoy effects in games. *Organizational Behavior and Human Decision Processes*, 104(2):193–206.

- [Conradie et al., 2006] Conradie, W., Goranko, V., and Vakarelov, D. (2006). Algorithmic correspondence and completeness in modal logic i: The core algorithm schema. *Logical Methods in Computer Science*, 2(1):1–26.
- [Cressman and Tao, 2014] Cressman, R. and Tao, Y. (2014). The replicator equation and other game dynamic. *Proc. Natl. Acad. Sci. U.S.A.*, pages 10810–10817. doi:10.1073/pnas.1400823111.
- [de Finetti, 1931] de Finetti, B. (1931). Sul Significato Soggettivo della Probabilità. *Fundamenta Mathematicae*, 17:298–329.
- [de Weerd et al., 2013] de Weerd, H., Verbrugge, R., and Verheij, B. (2013). How much does it help to know what she knows you know? an agent-based simulation study. *Artificial Intelligence*, 199-200:67–92. M1 - Journal Article.
- [Dekel et al., 2007] Dekel, E., Ely, J. C., and Ylankaya, O. (2007). Evolution of Preferences. *The Review of Economic Studies*, 74:3:685–704.
- [DeMiguel et al., 2009] DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies*, 22(5):1915–1953.
- [Dempster, 1967] Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Statist.*, 38(2):325–339.
- [Ellsberg, 1961] Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75:4:643–669.
- [Fagin, 1994] Fagin, R. (1994). A quantitative analysis of modal logic. *Journal of Symbolic Logic*, 59(1):209–252.
- [Fagin et al., 1999] Fagin, R., Geanakoplos, J., Halpern, J., and Vardi, M. (1999). The Hierarchical Approach to Modeling Knowledge and Common Knowledge. *International Journal of Game Theory*, 28(3):331–365.
- [Fagin et al., 1995] Fagin, R., Halpern, J., Moses, Y., and Vardi, M. (1995). *Reasoning about Knowledge*. MIT Press.
- [Fagin and Halpern, 1990] Fagin, R. and Halpern, J. Y. (1990). A New Approach to Updating Beliefs. In *Proceedings of the 6th Conference on Uncertainty in AI*, pages 317–325. New York: Elsevier.
- [Fagin et al., 1991] Fagin, R., Halpern, J. Y., and Vardi, M. Y. (1991). A model-theoretic analysis of knowledge. *Journal of the ACM*, 91(2):382–428.
- [Fawcett et al., 2013] Fawcett, T. W., Hamblin, S., and Giraldeaub, L.-A. (2013). Exposing the behavioral gambit: the evolution of learning and decision rules. *Behavioral Ecology*, 24(1):2–11.

- [Fehr and Schmidt, 1999] Fehr, E. and Schmidt, K. M. (1999). A Theory of Fairness, Competition and Cooperation. *Quarterly Journal of Economics*, 114:817–868.
- [Fishburn, 1982] Fishburn, P. C. (1982). Nontransitive measurable utility. *Journal of Mathematical Psychology*, 26(1):31 – 67.
- [Franke, 2012] Franke, M. (2012). On Assertoric and Directive Signals and the Evolution of Dynamic Meaning. *International Review of Pragmatics*, 4:232–260.
- [Friedman, 1953] Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press.
- [Fudenberg and Tirole, 1991] Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, Cambridge, MA.
- [Gardenfors and Sahlin, 1982] Gardenfors, P. and Sahlin, N.-E. (1982). Unreliable Probabilities, Risk Taking, and Decision Making. *Synthese*, 53:361–386.
- [Gigerenzer, 2008] Gigerenzer, G. (2008). Why heuristics work. *Perspectives on Psychological Science*, 3(1):20–29.
- [Gigerenzer and Brighton, 2009] Gigerenzer, G. and Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1):107–143.
- [Gigerenzer and Goldstein, 1996] Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669.
- [Gilboa, 2010] Gilboa, I. (2010). Questions in Decision Theory. *Annual Reviews in Economics*, 2:1–19.
- [Gilboa, 2014] Gilboa, I. (2014). Rationality and the Bayesian Paradigm: An Integrative Note. Unpublished manuscript.
- [Gilboa, 2015] Gilboa, I. (2015). Rationality and the Bayesian Paradigm. *Journal of Economic Methodology*, 22(3):312–334.
- [Gilboa et al., 2010] Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. (2010). Objective and subjective rationality in a multiple prior model. *Econometrica*, 78(2):755–770.
- [Gilboa and Marinacci, 2013] Gilboa, I. and Marinacci, M. (2013). Ambiguity and the bayesian paradigm. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics*, volume 1, pages 179–242. Cambridge University Press. Cambridge Books Online.

- [Gilboa et al., 2009] Gilboa, I., Postlewaite, A., and Schmeidler, D. (2009). Is it always rational to satisfy savage's axioms? *Economics and Philosophy*, 25(3):285–296.
- [Gilboa et al., 2012] Gilboa, I., Postlewaite, A., and Schmeidler, D. (2012). Rationality of belief or: Why savage's axioms are neither necessary nor sufficient for rationality. *Synthese*, 187(1):11–31.
- [Gilboa and Schmeidler, 1989] Gilboa, I. and Schmeidler, D. (1989). Maxmin Expected Utility with Non-unique Prior. *Journal of Mathematical Economics*, 18:141–153.
- [Gilboa and Schmeidler, 1995] Gilboa, I. and Schmeidler, D. (1995). Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639.
- [Gilboa and Schmeidler, 2001] Gilboa, I. and Schmeidler, D. (2001). *A Theory of Case-Based Decisions*. Cambridge University Press.
- [Hagen et al., 2012] Hagen, E. H., Chater, N., Gallistel, C. R., Houston, A., Kacelnik, A., Kalenscher, T., Nettle, D., Oppenheimer, D., and Stephens, D. W. (2012). Decision making: What can evolution do for us? In Hammerstein, P. and Stevens, J. R., editors, *Evolution and the Mechanisms of Decision Making*, pages 97–126. MIT Press, Cambridge, MA.
- [Halpern and Pass, 2012] Halpern, J. Y. and Pass, R. (2012). Iterated Regret Minimization: a New Solution Concept. *Games and Economic Behavior*, 74:1:194–207.
- [Harsanyi, 1967] Harsanyi, J. (1967). Games with Incomplete Information Played by “Bayesian” Players, Part I. The Basic Model. *Management Science*, 14(3):159–182.
- [Harsanyi, 1968a] Harsanyi, J. (1968a). Games with Incomplete Information Played by “Bayesian” Players, Part II. Bayesian Equilibrium Points. *Management Science*, 14(5):320–334.
- [Harsanyi, 1968b] Harsanyi, J. (1968b). Games with Incomplete Information Played by “Bayesian” Players, Part III. The Basic Probability Distribution of the Game. *Management Science*, 14(7):486–502.
- [Harsanyi and Selten, 1988] Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press.
- [Hayashi, 2008] Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1):242–268.

- [Heifetz et al., 2013] Heifetz, A., Meier, M., and Schipper, B. C. (2013). Dynamic unawareness and rationalizable behavior. *Games and Economic Behavior*, 81:50–68.
- [Heifetz and Mongin, 2001] Heifetz, A. and Mongin, P. (2001). Probability Logic for Type Spaces. *Games and Economic Behavior*, 35:31–53.
- [Heifetz and Samet, 1998] Heifetz, A. and Samet, D. (1998). Knowledge spaces with arbitrarily high rank. *Games and Economic Behavior*, 22:260–273.
- [Ho et al., 1998] Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests.". *American Economic Review*, 88(4):947–69.
- [Hofbauer, 1985] Hofbauer, J. (1985). The selection mutation equation. *Journal of Mathematical Biology*, 23(1):41–53.
- [Hofbauer and Sigmaund, 1998] Hofbauer, J. and Sigmaund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- [Huttegger and Zollman, 2010] Huttegger, S. M. and Zollman, K. J. S. (2010). Dynamic stability and basins of attraction in the Sir Philip Sidney game. *Proceedings of the Royal Society of London B: Biological Sciences*, 277(1689):1915–1922.
- [Huttegger and Zollman, 2013] Huttegger, S. M. and Zollman, K. J. S. (2013). Methodology in Biological Game Theory. *British Journal for the Philosophy of Science*, 64(3):637–658.
- [Huygens, 1657] Huygens, C. (1657). *De Ratiociniis in Ludo Aleae*. Amsterdam: van Schooten.
- [Jeffrey, 1990] Jeffrey, R. C. (1990). *The Logic of Decision*. University of Chicago Press.
- [Kahneman and Tversky, 1979] Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.
- [Kajii and Ui, 2005] Kajii, A. and Ui, T. (2005). Incomplete Information Games with Multiple Priors. *The Japanese Economic Review*, 56:332–351.
- [Kraft et al., 1959] Kraft, C. H., Pratt, J. W., and Seidenberg, A. (1959). Intuitive probability on finite sets. *The Annals of Mathematical Statistics*, 30(2):408–419.
- [Kraus and Lehmann, 1988] Kraus, S. and Lehmann, D. J. (1988). Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174.

- [Laplace, 1814] Laplace, P. S. (1814). *Théorie analytique des probabilités*. Paris: Ve. Courcier.
- [Lehmann, 1996] Lehmann, D. (1996). Generalized Qualitative Probability: Savage revisited. In *Proceeding of the 12th Annual Conference on Uncertainty in Artificial Intelligence*, pages 381–388. Morgan-Kaufmann, San Francisco.
- [Levi, 1974] Levi, I. (1974). On Indeterminate Probabilities. *The Journal of Philosophy*, 71:13:391–418.
- [Liu, 2015] Liu, H. (2015). *Integrating Game Theory and Modeling Uncertainty Using Imprecise Probabilities*. PhD thesis, Princeton University.
- [Loomes and Sugden, 1982] Loomes, G. and Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, 92:368:805–824.
- [Lorini, 2013] Lorini, E. (2013). On the epistemic foundation for iterated weak dominance: an analysis in a logic of individual and collective attitudes. *Journal of Philosophical Logic*, 42(6):863–904.
- [Lorini and Schwarzenrüber, 2010] Lorini, E. and Schwarzenrüber, F. (2010). A modal logic of epistemic games. *Games*, 1(4):478–526.
- [Machina, 1987] Machina, M. J. (1987). Choice under uncertainty: Problems solved and unsolved. *Journal of Economic Perspectives*, 1(1):121–154.
- [Markowitz, 1952] Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.
- [Mas-Colell et al., 1995] Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- [Maynard Smith, 1974] Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47(1):209 – 221.
- [Maynard Smith, 1982] Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.
- [Maynard Smith and Price, 1973] Maynard Smith, J. and Price, G. R. (1973). The Logic of Animal Conflict. *Nature*, 246:15–18.
- [Mengel, 2012] Mengel, F. (2012). Learning across games. *Games and Economic Behavior*, 74(2):601 – 619.
- [Mertens and Zamir, 1985] Mertens, J. F. and Zamir, S. (1985). Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29.

- [Milnor, 1954] Milnor, J. W. (1954). Games against nature. In Thrall, R. M., Coombs, C. H., and Davis, R. L., editors, *Decision Processes*, pages 49–59. Wiley, New York.
- [Mohlin, 2012] Mohlin, E. (2012). Evolution of theories of mind. *Games and Economic Behavior*, 75(1):299 – 318.
- [Nash, 1950] Nash, J. F. (1950). *Non-Cooperative Games*. PhD thesis, Princeton University.
- [von Neumann and Morgenstern, 1944] von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, New Jersey.
- [Nowak, 2006] Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press.
- [Oliver, 2003] Oliver, A. (2003). A quantitative and qualitative test of the allais paradox using health outcomes. *Journal of Economic Psychology*, 24(1):35 – 48.
- [Osborne and Rubinstein, 1994] Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- [Pascal, 1670] Pascal, B. (1670). *Pensées*.
- [Ramsey, 1926] Ramsey, F. P. (1926). Truth and Probability. In Braithwaite, R. B., editor, *The Foundation of Mathematics and Other Logical Essays (1931)*, pages 156–198. Kegan, Paul, Trench, Trubner and Co.
- [Rayo and Becker, 2007] Rayo, L. and Becker, G. S. (2007). Evolutionary Efficiency and Happiness. *Journal of Political Economy*, 115:302–337.
- [Robalino and Robson, 2016] Robalino, N. and Robson, A. (2016). The Evolution of Strategic Sophistication. *American Economic Review*, 106(4):1046–72.
- [Robson and Samuelson, 2011] Robson, A. and Samuelson, L. (2011). The Evolutionary Foundations of Preferences. In Benhabib, J., Jackson, M., and Bisin, A., editors, *Handbook of Social Economics, 1A*. North-Holland.
- [Rubinstein, 2006] Rubinstein, A. (2006). Dilemmas of an economic theorist. *Econometrica*, 74(4):865–883.
- [Sahlqvist, 1975] Sahlqvist, H. (1975). Completeness and correspondence in the first and second order semantics for modal logics. In *Proceedings 3rd Scandinavian Logic Symposium 1973*, volume number 82 in Studies in Logic.

- [Samuelson, 1997] Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press.
- [Samuelson, 2001] Samuelson, L. (2001). Introduction to the Evolution of Preferences. *Journal of Economic Theory*, 97:225–230.
- [Sandholm, 2010] Sandholm, W. H. (2010). *Population games and evolutionary dynamics*. Economic learning and social evolution. Cambridge, Mass. MIT Press.
- [Savage, 1951] Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46:55–67.
- [Savage, 1954] Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.
- [Schmeidler, 1989] Schmeidler, D. (1989). Subjective Probability and Expected Utility Without Additivity. *Econometrica*, 57:571–589.
- [Scott, 1964] Scott, D. (1964). Measurement structures and linear inequalities. *Journal of Mathematical Psychology*, 1(2):233 – 247.
- [Shafer, 1976] Shafer, G. (1976). *A Mathematical Theory of Evidence*. Princeton University Press.
- [Shafir et al., 1993] Shafir, E., Simonson, I., and Tversky, A. (1993). Reason-based choice. *Cognition*, (49):11–36.
- [Shah and Oppenheimer, 2011] Shah, A. K. and Oppenheimer, D. M. (2011). Heuristics made easy: An effort-reduction framework. *Journal of Economic Theory*, 146:2226–2251.
- [Simon, 1955] Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69:99–118.
- [Simon, 1990] Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41(1):1–20. PMID: 18331187.
- [Simon, 1992] Simon, H. A. (1992). *Economics, bounded rationality, and the cognitive revolution*. Aldershot Hants, England: Elgar.
- [Sinervo and Lively, 1996] Sinervo, B. and Lively, C. M. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 380:240–243.
- [Skyrms, 1996] Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge University Press.

- [Skyrms, 2010] Skyrms, B. (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- [Spohn, 2012] Spohn, W. (2012). *The Laws of Belief. Ranking Theory and its Philosophical Applications*. Oxford University Press.
- [Stahl, 1993] Stahl, D. O. (1993). Evolution of Smart-n Players. *Games and Economic Behavior*, 5(4):604 – 617.
- [Stahl and Wilson, 1994] Stahl, D. O. and Wilson, P. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327.
- [Stahl and Wilson, 1995] Stahl, D. O. and Wilson, P. (1995). On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254.
- [Stenning and van Lambalgen, 2008] Stenning, K. and van Lambalgen, M. (2008). *Human Reasoning and Cognitive Science*. MIT Press.
- [Stoye, 2011] Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146:2226–2251.
- [Sugden, 1993] Sugden, R. (1993). An axiomatic foundation for regret theory. *Journal of Economic Theory*, 60(1):159 – 180.
- [Tan and Pearl, 1994] Tan, S.-W. and Pearl, J. (1994). Qualitative Decision Theory. In *Proceeding of the 12th National Conference on Artificial Intelligence*, pages 928–932. Morgan-Kaufmann, San Francisco.
- [Taylor and Jonker, 1978] Taylor, P. D. and Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Bioscience*, 40(1–2):145–156.
- [Trautmann and van de Kuilen, 2016] Trautmann, S. T. and van de Kuilen, G. (2016). Ambiguity Attitudes. In Keren, G. and Wu, G., editors, *The Wiley Blackwell Handbook of Judgment and Decision Making*. Wiley-Blackwell.
- [Tversky and Kahneman, 1974] Tversky, A. and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157):1124–1131.
- [Tversky and Kahneman, 1981] Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458.
- [Tversky and Simonson, 1993] Tversky, A. and Simonson, I. (1993). Context-dependent Preferences. *Manage. Sci.*, 39(10):1179–1189.

- [Wakker, 2010] Wakker, P. (2010). *Prospect Theory*. Cambridge University Press.
- [Wald, 1950] Wald, A. (1950). *Statistical Decision Functions*. John Wiley and Sons.
- [Walley, 1991] Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman and Hall.
- [Weibull, 1995] Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press.
- [Zollman, 2008] Zollman, K. J. S. (2008). Explaining fairness in complex environments. *Politics, Philosophy and Economics*, 7(1):81–97.
- [Zvesper, 2010] Zvesper, J. (2010). *Playing with Information*. PhD thesis, University of Amsterdam.

Abstract

This thesis is centered on the issue of *rational choice*. Traditional decision-theoretic arguments aim at providing axiomatizations of different norms for decision making, and evaluate their rationality on the basis of the normative strength of their axiom systems.

The approach taken in this dissertation is different. Instead of arguing in favor or against the normativity of some system of axioms, we decided to take a less well-trodden path: the study of the *ecological* rationality of different decision criteria.

To this end, we extended the standard single-game models used in evolutionary game theory to include a multitude of different interactive decision problems. We consider the introduction of such a *multi-game* model, called *the game of life*, a principal contribution of this work in itself, in that it allows to lift the focus of the investigation from simple behavior to general behavior-generating mechanisms.

The main results of this thesis concern the evolutionary competition between different ways of making choices in rich and complex environments. Classic decision criteria are compared from an ecological point of view, with respect to their evolutionary *fitness*, and regret-based principles prove to be especially beneficial in many (interactive) decision contexts.

Samenvatting

Dit proefschrift richt zich op *rationele keuze*. Traditionele argumenten in de besluitvormingstheorie streven ernaar om verschillende normen voor besluitvorming te axiomatiseren, en evalueren de rationaliteit hiervan op basis van de normatieve kracht van de resulterende axiomasystemen.

In dit proefschrift gebruiken we een andere aanpak. We argumenteren niet voor of tegen de normativiteit van een bepaald axiomasysteem. In plaats daarvan kiezen we voor een minder vaak gebruikte aanpak: het bestuderen van de *ecologische* rationaliteit van verschillende beslissingcriteria.

Om dit te doen hebben we de gebruikelijke single-game-modellen, die gebruikt worden in de evolutionaire speltheorie, uitgebreid zodat ze een groot aantal verschillende interactieve beslissingsproblemen omvatten. Het introduceren van een dergelijk multi-game-model, genaamd *the game of life*, beschouwen we als een belangrijke bijdrage van dit werk, omdat dit het mogelijk maakt om de focus van het onderzoek te verleggen van eenvoudig gedrag naar algemene mechanismen die gedrag voortbrengen.

De voornaamste resultaten van dit proefschrift gaan over de evolutionaire competitie tussen verschillende manieren om keuzes te maken in rijke en complexe omgevingen. We vergelijken klassieke beslissingscriteria vanuit een ecologisch standpunt, aan de hand van hun evolutionaire *fitness*. Principes die gebaseerd zijn op spijt blijken in het bijzonder van pas te komen in veel (interactieve) beslissingscontexten.

Titles in the ILLC Dissertation Series:

- ILLC DS-2009-01: **Jakub Szymanik**
Quantifiers in TIME and SPACE. Computational Complexity of Generalized Quantifiers in Natural Language
- ILLC DS-2009-02: **Hartmut Fitz**
Neural Syntax
- ILLC DS-2009-03: **Brian Thomas Semmes**
A Game for the Borel Functions
- ILLC DS-2009-04: **Sara L. Uckelman**
Modalities in Medieval Logic
- ILLC DS-2009-05: **Andreas Witzel**
Knowledge and Games: Theory and Implementation
- ILLC DS-2009-06: **Chantal Bax**
Subjectivity after Wittgenstein. Wittgenstein's embodied and embedded subject and the debate about the death of man.
- ILLC DS-2009-07: **Kata Balogh**
Theme with Variations. A Context-based Analysis of Focus
- ILLC DS-2009-08: **Tomohiro Hoshi**
Epistemic Dynamics and Protocol Information
- ILLC DS-2009-09: **Olivia Ladinig**
Temporal expectations and their violations
- ILLC DS-2009-10: **Tikitu de Jager**
"Now that you mention it, I wonder...": Awareness, Attention, Assumption
- ILLC DS-2009-11: **Michael Franke**
Signal to Act: Game Theory in Pragmatics
- ILLC DS-2009-12: **Joel Uckelman**
More Than the Sum of Its Parts: Compact Preference Representation Over Combinatorial Domains
- ILLC DS-2009-13: **Stefan Bold**
Cardinals as Ultrapowers. A Canonical Measure Analysis under the Axiom of Determinacy.
- ILLC DS-2010-01: **Reut Tsarfaty**
Relational-Realizational Parsing

- ILLC DS-2010-02: **Jonathan Zvesper**
Playing with Information
- ILLC DS-2010-03: **Cédric Dégrement**
The Temporal Mind. Observations on the logic of belief change in interactive systems
- ILLC DS-2010-04: **Daisuke Ikegami**
Games in Set Theory and Logic
- ILLC DS-2010-05: **Jarmo Kontinen**
Coherence and Complexity in Fragments of Dependence Logic
- ILLC DS-2010-06: **Yanjing Wang**
Epistemic Modelling and Protocol Dynamics
- ILLC DS-2010-07: **Marc Staudacher**
Use theories of meaning between conventions and social norms
- ILLC DS-2010-08: **Amélie Gheerbrant**
Fixed-Point Logics on Trees
- ILLC DS-2010-09: **Gaëlle Fontaine**
Modal Fixpoint Logic: Some Model Theoretic Questions
- ILLC DS-2010-10: **Jacob Vosmaer**
Logic, Algebra and Topology. Investigations into canonical extensions, duality theory and point-free topology.
- ILLC DS-2010-11: **Nina Gierasimczuk**
Knowing One's Limits. Logical Analysis of Inductive Inference
- ILLC DS-2010-12: **Martin Mose Bentzen**
Stit, It, and Deontic Logic for Action Types
- ILLC DS-2011-01: **Wouter M. Koolen**
Combining Strategies Efficiently: High-Quality Decisions from Conflicting Advice
- ILLC DS-2011-02: **Fernando Raymundo Velazquez-Quesada**
Small steps in dynamics of information
- ILLC DS-2011-03: **Marijn Koolen**
The Meaning of Structure: the Value of Link Evidence for Information Retrieval
- ILLC DS-2011-04: **Junte Zhang**
System Evaluation of Archival Description and Access

- ILLC DS-2011-05: **Lauri Keskinen**
Characterizing All Models in Infinite Cardinalities
- ILLC DS-2011-06: **Rianne Kaptein**
Effective Focused Retrieval by Exploiting Query Context and Document Structure
- ILLC DS-2011-07: **Jop Briët**
Grothendieck Inequalities, Nonlocal Games and Optimization
- ILLC DS-2011-08: **Stefan Minica**
Dynamic Logic of Questions
- ILLC DS-2011-09: **Raul Andres Leal**
Modalities Through the Looking Glass: A study on coalgebraic modal logic and their applications
- ILLC DS-2011-10: **Lena Kurzen**
Complexity in Interaction
- ILLC DS-2011-11: **Gideon Borensztajn**
The neural basis of structure in language
- ILLC DS-2012-01: **Federico Sangati**
Decomposing and Regenerating Syntactic Trees
- ILLC DS-2012-02: **Markos Mylonakis**
Learning the Latent Structure of Translation
- ILLC DS-2012-03: **Edgar José Andrade Lotero**
Models of Language: Towards a practice-based account of information in natural language
- ILLC DS-2012-04: **Yurii Khomskii**
Regularity Properties and Definability in the Real Number Continuum: idealized forcing, polarized partitions, Hausdorff gaps and mad families in the projective hierarchy.
- ILLC DS-2012-05: **David García Soriano**
Query-Efficient Computation in Property Testing and Learning Theory
- ILLC DS-2012-06: **Dimitris Gakis**
Contextual Metaphilosophy - The Case of Wittgenstein
- ILLC DS-2012-07: **Pietro Galliani**
The Dynamics of Imperfect Information

- ILLC DS-2012-08: **Umberto Grandi**
Binary Aggregation with Integrity Constraints
- ILLC DS-2012-09: **Wesley Halcrow Holliday**
Knowing What Follows: Epistemic Closure and Epistemic Logic
- ILLC DS-2012-10: **Jeremy Meyers**
Locations, Bodies, and Sets: A model theoretic investigation into nominalistic mereologies
- ILLC DS-2012-11: **Floor Sietsma**
Logics of Communication and Knowledge
- ILLC DS-2012-12: **Joris Dormans**
Engineering emergence: applied theory for game design
- ILLC DS-2013-01: **Simon Pauw**
Size Matters: Grounding Quantifiers in Spatial Perception
- ILLC DS-2013-02: **Virginie Fiutek**
Playing with Knowledge and Belief
- ILLC DS-2013-03: **Giannicola Scarpa**
Quantum entanglement in non-local games, graph parameters and zero-error information theory
- ILLC DS-2014-01: **Machiel Keestra**
Sculpting the Space of Actions. Explaining Human Action by Integrating Intentions and Mechanisms
- ILLC DS-2014-02: **Thomas Icard**
The Algorithmic Mind: A Study of Inference in Action
- ILLC DS-2014-03: **Harald A. Bastiaanse**
Very, Many, Small, Penguins
- ILLC DS-2014-04: **Ben Rodenhäuser**
A Matter of Trust: Dynamic Attitudes in Epistemic Logic
- ILLC DS-2015-01: **María Inés Crespo**
Affecting Meaning. Subjectivity and evaluativity in gradable adjectives.
- ILLC DS-2015-02: **Mathias Winther Madsen**
The Kid, the Clerk, and the Gambler - Critical Studies in Statistics and Cognitive Science

- ILLC DS-2015-03: **Shengyang Zhong**
Orthogonality and Quantum Geometry: Towards a Relational Reconstruction of Quantum Theory
- ILLC DS-2015-04: **Sumit Sourabh**
Correspondence and Canonicity in Non-Classical Logic
- ILLC DS-2015-05: **Facundo Carreiro**
Fragments of Fixpoint Logics: Automata and Expressiveness
- ILLC DS-2016-01: **Ivano A. Ciardelli**
Questions in Logic
- ILLC DS-2016-02: **Zoé Christoff**
Dynamic Logics of Networks: Information Flow and the Spread of Opinion
- ILLC DS-2016-03: **Fleur Leonie Bouwer**
What do we need to hear a beat? The influence of attention, musical abilities, and accents on the perception of metrical rhythm
- ILLC DS-2016-04: **Johannes Marti**
Interpreting Linguistic Behavior with Possible World Models
- ILLC DS-2016-05: **Phong Lê**
Learning Vector Representations for Sentences - The Recursive Deep Learning Approach
- ILLC DS-2016-06: **Gideon Maillette de Buy Wenniger**
Aligning the Foundations of Hierarchical Statistical Machine Translation
- ILLC DS-2016-07: **Andreas van Cranenburgh**
Rich Statistical Parsing and Literary Language
- ILLC DS-2016-08: **Florian Speelman**
Position-based Quantum Cryptography and Catalytic Computation
- ILLC DS-2016-09: **Teresa Piovesan**
Quantum entanglement: insights via graph parameters and conic optimization
- ILLC DS-2016-10: **Paula Henk**
Nonstandard Provability for Peano Arithmetic. A Modal Perspective



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION