# Acknowledgments

Writing this book has been a long and adventurous journey on which I embarked five years ago, in 2018. Now that I have completed the journey I can say that I enjoyed every moment of it, and I hope that the reader will sense this while reading the book. I would now like to take the opportunity to thank a few people that have been fundamental for writing this book.

Important parts of the book are based on research that I have carried out together with my co-authors Christian Bach and Stephan Jagau. With Christian Bach I have developed a line of research in which we explore the reasoning foundations of games with incomplete information, which has contributed to parts of the Chapters 5 and 6. It has always been a huge pleasure to work together with you, Christian, and to have you come over at my office in Maastricht for some nice research sessions at the whiteboard, followed by a lunch at our favorite café in Maastricht. I hope we can continue our collaboration for many years to come. And, importantly, thanks for being such a good friend.

Some years ago, Stephan Jagau and I have started working on the reasoning foundations for psychological games, leading to insights that are discussed in Chapter 8. Recently, our collaboration mainly took place via zoom, because Stephan was based in California and I am based in the Netherlands. I cherish these sessions tremendously, and I hope we can expand our collaboration during the next few years, now that Stephan is based in Nottingham – a bit closer to Maastricht. I would also like to thank Stephan for our discussions on one-person decision problems that helped to shape Chapter 2.

Parts of the Chapters 8 and 9 build upon the work of my former PhD student Niels Mourmans, who devoted his PhD dissertation to the theoretical foundations of psychological games. I can still remember the many fruitful and enjoyable sessions we had in my office, together with the other supervisor Elias Tsakas. It resulted in a very nice PhD dissertation.

I would also like to thank Yumo Chen, Li Junkang, Shuige Liu and, again, Niels Mourmans, who read parts of the book and provided me with some highly appreciated feedback.

In 2023 I have given a two day mini-course based on this book at the Universidad Carlos III de Madrid and at my own Maastricht University. I am grateful to the audiences at these mini-courses for their feedback and comments.

An important part of this book has been written while I was visiting the Department of Economics at the Universidad Carlos III de Madrid, in 2022 and 2023. I would like to thank Universidad Carlos III de Madrid for their hospitality. My former colleagues at this university, and the wonderful city of Madrid, have been a strong source of inspiration for the book.

Right from the beginning, the collaboration with Cambridge University Press has been very smooth and pleasant. I would like to send a special word of gratitude to the reviewers for their excellent feedback on my manuscript and proposal, and to Philip Good for his excellent help during the past year.

Last but not least, I would like to pay a tribute to my friends, my brother and sister, my parents, my children and my wife for their continuous mental support, which filled the journey of writing this book with warm and sunny days.

# Chapter 1

# Introduction

**Why this book?** In a game you must make a decision, knowing that the final consequence also depends on the decisions of others. The preferences over your own decisions will therefore crucially depend on what you believe that these others will do. In order to make a good decision you must reason about the decisions of others, form a belief about the likely decisions that they will make, and finally base your own decision on this belief.

In many situations of interest, you are either not fully informed about the preferences of the other decision makers, resulting in a game with *incomplete information,* or you are unaware of some of the decisions that the others can make, yielding a game with *unawareness,* or the preferences over your own decisions may depend on what you believe that some other person believes that you will do, giving rise to a *psychological* game.

Each of these classes of games has been studied intensively in the game theoretic literature. Yet, many of the papers involved may pose a serious challenge for the reader in terms of complexity and technicality. In addition, the three strands of literature that investigate the three classes of games above are so different that the reader may get the impression of entering three completely separate worlds.

One purpose of this book is to explore these three classes of games in an accessible and unified manner. Indeed, we try to keep the technicalities in this book at a minimum while still remaining fully rigorous. The many examples that illustrate the concepts in the book will help the reader to understand the various pieces of theory being discussed. Moreover, the three classes of games will be modelled as collections of one-person decision problems under uncertainty – one for each player – thereby providing a unified decision theoretic approach. Also, for each of these classes of games we focus on the reasoning concepts of *common belief in rationality,* common belief in rationality with *correct beliefs*, and common belief in rationality with *symmetric beliefs,* leading to a unified reasoning approach. This unified approach will hopefully demonstrate that, despite their differences, the three classes of games can still be analyzed in essentially the same fashion.

Another goal of this book is to establish a natural bridge between one-person decision theory on the one hand and game theory on the other hand. If one compares the literature on one-person decision

theory with that of game theory it is again striking how different these are. It often seems that one enters a completely new world when moving from one-person decision theory to game theory. This book attempts to provide a smooth and natural transition between one-person decision theory and game theory, by modelling games as collections of one-person decision problems – one for every player.

**From decision theory to game theory.** In a one-person decision problem under uncertainty, the decision maker must reach a decision while knowing that the final consequence will also depend on uncertain events that are beyond his control. Such uncertain events are usually called *states.* Every player in a game therefore faces a particular one-person decision problem under uncertainty where the states concern the possible choices that can be made by his opponents. In a psychological game, where the player's preferences over his own choices may depend on what another player believes about his choices, the states even contain the possible *beliefs* that another player can have about his own choices.

In this book we indeed model a game as a collection of one-person decision problems – one for every player in the game. In Part I of this book we start by discussing one-person decision problems under uncertainty, whereas the subsequent chapters show how the various classes of games discussed above can all be identified with particular one-person decision problems. By doing so we thus provide a smooth transition from one-person decision theory to game theory. Moreover, it is shown that the various classes of games, despite their differences, can all be modelled in a similar fashion as collections of one-person decision problems, thereby offering a unified decision theoretic approach to these classes of games.

**Reasoning about decision problems of others.** If we model a game as a collection of decision problems, as we do in this book, it is important for a player to actively reason about the decision problems of others. Indeed, the states in the decision problem of a player involve the possible choices that other players can make. Some opponents' choices are more plausible than others, and hence some states in the decision problem seem more probable than others. To form a good belief about the states, a player must therefore reason about the likely choices that his opponents will make, which, in turn, forces him to reason about the opponents' decision problems. It is this interactive element that distinguishes game theory from "neutral" one-person decision theory, where the states are determined by nature and not by the choices of human beings.

**Unified reasoning approach.** In Parts II, III, IV and V we explore standard games, where the preferences and the sets of choices are commonly known among the players, games with incomplete information, games with unawareness and psychological games, respectively. For each of those classes we concentrate on three prominent lines of reasoning: (i) *common belief in rationality,* stating that a player believes that all others choose rationally, believes that all others believe that all others choose rationally, and so on, (ii) common belief in rationality in combination with a *simple belief hierarchy,* where a player believes that the others are correct about the beliefs he holds, and (iii) common belief in rationality in combination with a *symmetric belief hierarchy,* where a player believes that his beliefs about others are, in a sense, "symmetric" to the beliefs he expects others to hold about himself.

It turns out that these three reasoning concepts can be defined in a very similar way for each of the four classes of games mentioned above. We thus provide a unified reasoning approach to the different classes of games discussed in this book. It also shows the robustness of *epistemic game theory* – the approach to game theory that explores the reasoning of players before they make a choice. Indeed, the epistemic conditions on the players' beliefs that are needed to formally define the three reasoning concepts above are nearly identical across the various classes of games we consider.

| | common belief in rationality with ... | | |
|---|---|---|---|
| | ... | simple belief hierarchy | symmetric belief hierarchy |
| standard game | iterated elimination of strictly dominated choices | Nash equilibrium | correlated equilibrium |
| incomplete information | generalized iterated strict dominance | generalized Nash equilibrium | Bayesian equilibrium |
| unawareness | iterated strict dominance for unawareness | | |
| psychological game | iterated elimination of choices and second-order expectations | psychological Nash equilibrium | psychological correlated equilibrium |

Table 1.1 Reasoning concepts and their associated solution concepts

**From reasoning to choices.** For each of the three reasoning concepts above we are eventually interested in the *choices* that the players can rationally make if they reason along those lines. Mappings that select for every game a set of "plausible" choices for every player are usually called *solution concepts.* It is shown in this book that for every class of games, each of the three reasoning concepts above gives rise to a solution concept that is either well-known, not so well-known, or perhaps even new.

In Part II, for standard games, we argue that the first reasoning concept above can be characterized by the solution concept known as *iterated elimination of strictly dominated choices,* that the second reasoning concept corresponds to *Nash equilibrium,* whereas the third reasoning concept leads to *correlated equilibrium.* All of these are well-known solution concepts in game theory.

In the subsequent parts we explain that for the other three classes of games, the three reasoning concepts above lead to appropriate variations of the iterated elimination of strictly dominated choices, Nash equilibrium and correlated equilibrium, respectively. It turns out, for instance, that for games with incomplete information, common belief in rationality in combination with a symmetric belief hierarchy leads to the famous solution concept of *Bayesian equilibrium,* showing that Bayesian equilibrium is the incomplete information counterpart to correlated equilibrium. In Table 1.1 we have indicated, for each of the four classes of games, and each of the three reasoning concepts, the resulting solution concept.

Note that the last two cells corresponding to games with unawareness are empty. The reason, as we will explain in Chapter 7, is that simple and symmetric belief hierarchies necessarily lead us to trivial cases of unawareness, where a player believes that all others are aware of the same sets of choices as he is aware of himself. This, however, would move us back to a standard game, and hence no new insights can be gained here.

The table shows that for psychological games, the conditions of common belief in rationality together with a symmetric belief hierarchy correspond to the solution concept of psychological correlated equilibrium – a concept that, to the best of our knowledge, is new. The other solution concepts in the table are taken from the existing literature.

**Reasoning first.** In every game-theoretic chapter we start by informally describing a line of *reasoning*, after which we formulate the reasoning concept rigorously by means of epistemic conditions within a mathematical model. Only afterwards will we search for a solution concept that characterizes those choices that a player can rationally make if he reasons in line with this concept. We believe this is the right order of doing things, as a player first reasons about the game, and only afterwards makes a choice based on this reasoning.

**One-person perspective.** In this book we take a *one-person perspective* to analyzing games. That is, we always put ourselves in the shoes of a single player, and explore how this player would reason about his opponents before making a choice. By taking this one-person perspective we bring the analysis of a game closer to how we would explore a one-person decision problem under uncertainty. Also this will help to establish a smooth transition from one-person decision theory to game theory.

**Examples from everyday life.** On purpose, we have selected examples from everyday life to illustrate the concepts in this book. The reason is that the book aims at readers from a large spectrum of disciplines, including economics, philosophy, logic, computer science, artificial intelligence, mathematics, political science and other related areas. Therefore, selecting examples from economics or politics, for instance, would perhaps not appeal so much to readers from other fields. Instead, we hope that every reader can identify with the examples from everyday life.

The main character in each of the examples is the reader, addressed as "you", who is facing game-theoretic scenarios together with his imaginary friends Barbara and Chris. The main question in these examples is always: What would you do in this situation, and why? We then use the reasoning concept of that chapter, or section, to answer this question.

**In-chapter questions and problems.** In the main text of every chapter there are short questions, referred to as "in-chapter questions", which will serve as a test for the reader to see whether he understands the concepts that are currently being explained. If the reader has difficulties answering a particular question, then this is probably a good reason to study the associated text once more.

The problems at the end of every chapter have a different purpose: They help the reader to apply the concepts of that chapter to analyze a new game-theoretic scenario from everyday life. The main characters in every problem are again you and Barbara, possibly together with Chris, in case there is a third player in the game. Every problem always starts with a story, and the first question of each problem is to translate this story into a game of a particular class. We find this translation step very important.

**Overview of related literature to every chapter.** At the end of each chapter we provide an overview of some of the papers that are related to the concepts which have been discussed in that chapter. Since this is a textbook, we do not give such references in the main text as to not distract the reader too much from the content.

**Restriction to static games.** In this book we restrict to *static* games, in which players only choose once, and receive no information about the choices of their opponents when making their own choice. As a consequence, we do not explore *dynamic* games in which a player, before making his choice, may completely or partially observe what other players before him have chosen. We do so because otherwise the book would become too big. In any case, the purpose of this book is not to give a

complete overview of the literature on games with incomplete information, games with unawareness, and psychological games, but rather to provide a unified treatment of these classes of games by focusing on three reasoning concepts for static games: common belief in rationality, common belief in rationality in combination with a simple belief hierarchy, and common belief in rationality in combination with a symmetric belief hierarchy. The literature offers, in addition, an extensive analysis of dynamic games for each of these three classes of games. But again, this topic falls outside the scope of this book.

**Economic applications.** There is an online appendix to this book which contains economic applications for each of the chapters. These chapters will be particularly interesting for readers with a background in economics, or teachers who use this book for a course within an economics or business program. The economic applications for the game theory chapters – that is, all chapters except Chapter 2 – focus on the Bertrand competition model where two firms compete in prices, and the Cournot competition model where two firms compete in quantities. It is shown, for instance, how such models can be extended to scenarios where there is incomplete information, unawareness, and where the preferences of the firms give rise to a psychological game.

**Axiomatic characterization of expected utility.** In Chapter 2 of the online appendix we provide an axiomatic characterization of expected utility for the model of one-person decision problems we use in this book. We have moved this part to the online appendix because it is rather advanced, and not essential for a good understanding of the remainder of the book. However, the reader who is interested in the assumptions behind the expected utility paradigm is very welcome to study this part of the online appendix.

**How to use the book?** This book can be used for self-study, but is also perfectly suited as a textbook for a course on game theory within a variety of study programs, including economics, mathematics, logic, philosophy, computer science, artificial intelligence, political sciences, and related areas. Moreover, the book can be used on different levels, such as the second or third year of a bachelor program, on a master level, or at a PhD level. The book provides enough flexibility here, because it contains many examples, gives verbal, intuitive descriptions for each of the concepts being discussed, but also rigorous mathematical descriptions of these concepts. Also, the more difficult parts in the book carry an asterisk. Depending on the level, the teacher may want to focus more on the intuitive descriptions and leave out the parts with an asterisk, or dive deeper into the rigorous mathematical descriptions and include the parts with an asterisk. On a PhD level, the teacher may also wish to discuss some of the proofs which can always be found at the end of every chapter. The examples, and the problems at the end of the chapters, however, remain essential for *every* level on which this course may be given.

The teacher may either use all chapters of the book, or a selection of those, for a course. All chapters would be suitable for an eight week course, or longer, where every week consists of one or two theory lectures, and an exercise session where the problems at the end of the chapters will be discussed. One could also opt for a course that covers one-person decision problems (Chapter 2), standard games (Chapters 3 and 4), and games with incomplete information (Chapters 5 and 6). Such a course would be appropriate for a bachelor program, for instance. Or one could completely focus on common belief in rationality, and leave out the parts on simple and symmetric belief hierarchies. Such a course would cover the Chapters 2, 3, 5, 7 and 8, with the possibility of leaving out Chapters 7 and 8. This would be ideal for a short mini-course. Apart from these, many other selections are possible here. However, every course should at least contain the part on one-person decision problems (Chapter 2 ) and the part on common belief in rationality in standard games (Chapter 3), as these constitute the basis for all further chapters.

**Related books, chapters and overview papers.** What is special about this book is that it zooms in on the reasoning of players before they make a choice in a game. It thereby takes an *epistemic game theory* perspective, by imposing conditions on the belief hierarchies of players, and subsequently characterizing those choices that can rationally be made under these conditions. Here, a belief hierarchy describes what a player believes about his opponents' choices, what he believes about the beliefs that his opponents have about the choices of others, and so on. It can be viewed as the end-product of a particular reasoning process.

Other books that take an epistemic game theory approach are the textbook by Perea (2012), the book in progress by Battigalli, Friedenberg and Siniscalchi (2023) and the monograph by Asheim (2006). The present book can be viewed as a follow-up on the book Perea (2012), which does not cover games with incomplete information, games with unawareness or psychological games, but does analyze dynamic games, as well as cautious reasoning modelled by lexicographic beliefs. In a sense, the present book is complementary to Perea (2012), and can be read perfectly without having read Perea (2012) before. Even more, the present book is completely self-contained, and can be read without any pre-knowledge about decision theory or game theory.

For an overview of the epistemic game theory approach one can also consult the handbook chapters by Dekel and Siniscalchi (2015), Bonanno (2015) and Perea (2021), the chapter by Pacuit and Roy (2017) and the survey papers by Brandenburger (1992, 2007), Geanakoplos (1992), Dekel and Gul (1997), Battigalli and Bonanno (1999), Board (2002) and Perea (2014).